

RANP: Resource Aware Neuron Pruning at Initialization for 3D CNNs

Zhiwei Xu^{1,3} · Thalaiyasingam Ajanthan^{1,4} · Vibhav Vineet² · Richard Hartley¹

Received: date / Accepted: date

Abstract Although 3D Convolutional Neural Networks are essential for most learning based applications involving dense 3D data, their applicability is limited due to excessive memory and computational requirements. Compressing such networks by pruning therefore becomes highly desirable. However, pruning 3D CNNs is largely unexplored possibly because of the complex nature of typical pruning algorithms that embeds pruning into an iterative optimization paradigm. In this work, we introduce a Resource Aware Neuron Pruning (RANP) algorithm that prunes 3D CNNs at initialization to high sparsity levels. Specifically, the core idea is to obtain an importance score for each neuron based on their sensitivity to the loss function. This neuron importance is then reweighted according to the neuron resource consumption related to FLOPs or memory. We demonstrate the effectiveness of our pruning method on 3D semantic segmentation with widely used 3D-UNets on ShapeNet and BraTS'18 datasets, video classification with MobileNetV2 and I3D on UCF101 dataset, and two-view stereo matching with Pyramid Stereo Matching (PSM) network on SceneFlow dataset. In these experiments, our RANP leads to roughly 50%-95%

Zhiwei Xu
zhiwei.xu@anu.edu.au

Thalaiyasingam Ajanthan
thalaiyasingam.ajanthan@anu.edu.au

Vibhav Vineet
vibhav.vineet@microsoft.com

Richard Hartley
richard.hartley@anu.edu.au

¹ Australian National University, Canberra, Australia
Australian Centre for Robotic Vision, Canberra, Australia
² Microsoft Research, Redmond, USA
³ Data61, CSIRO, Canberra, Australia
⁴ Amazon, Adelaide, Australia (work done prior to joining Amazon)

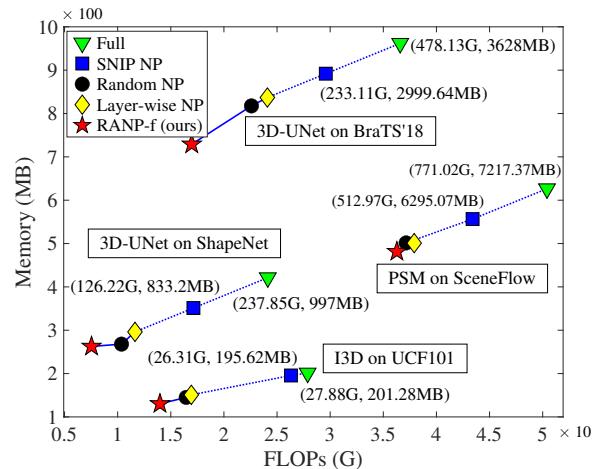


Fig. 1: **Bottom-left is the best.** Comparison of neuron pruning methods. Values of “PSM on SceneFlow” are divided by 10 for visualization. “Full” and “SNIP NP” are not drawn by scale with FLOPs (G) and memory (MB) values next to the markers. Our RANP-f performs best with large resource reductions while maintaining accuracy. More details of resource consumption with accuracy are provided in Table 2.

reduction in FLOPs and 35%-80% reduction in memory with negligible loss in accuracy compared to the unpruned networks. This significantly reduces the computational resources required to train 3D CNNs. The pruned network obtained by our algorithm can also be easily scaled up and transferred to another dataset for training.

Keywords Neuron pruning · Resource efficient · 3D CNNs · 3D semantic segmentation · Video classification · Stereo vision

1 Introduction

3D image analysis is important in various real-world applications including scene understanding Zhang et al. (2014, 2018), object recognition Ji et al. (2013); Hou et al. (2019), medical image analysis Cicek et al. (2016); Zanjani et al. (2019); Kleesiek et al. (2016), and video action recognition Karpathy et al. (2014); Simonyan and Zisserman (2014). Typically, sparse 3D data can be represented using point clouds Yi et al. (2016) whereas volumetric representation is required for dense 3D data which arises in domains such as medical imaging Menze et al. (2015) and video segmentation and classification Zhang et al. (2014); Karpathy et al. (2014); Simonyan and Zisserman (2014). While efficient neural network architectures can be designed for sparse point cloud data Graham et al. (2018); Qi et al. (2017), conventional dense 3D Convolutional Neural Network (CNN) is required for volumetric data. Such 3D CNNs are computationally expensive with excessive memory requirements for large-scale 3D tasks. Therefore, it is highly desirable to reduce the memory and FLOPs required to train 3D CNNs while maintaining the accuracy. This will not only enable large-scale applications but also 3D CNN training on resource-limited devices.

Network pruning is a prominent approach to compress a neural network by reducing the number of parameters or the number of neurons in each layer Guo et al. (2016); Dong et al. (2017); He et al. (2017); Han et al. (2015). However, most of the network pruning methods aim at 2D CNNs while pruning 3D CNNs is largely unexplored. This is mainly because pruning is typically targeted at reducing the test-time resource requirements while computational requirements of training time are as large as (if not more than) the unpruned network. Such pruning schemes are not suitable for 3D CNNs with dense volumetric data where training-time resource requirement is prohibitively large.

In this work, we introduce a Resource¹ Aware Neuron Pruning (RANP) that *prunes 3D CNNs at initialization*. Our method is inspired by, but superior to, SNIP Lee et al. (2019) which prunes redundant parameters of a network at initialization and only tests with small scale 2D CNNs for image classification. With the same characteristics of effectively pruning at initialization without requiring large computational resources, RANP yields better-pruned networks compared to SNIP by removing neurons that largely contribute to the high resource requirement. In our experiments on video classification and more challenging 3D semantic segmentation, with minimal accuracy loss, RANP yields 50%-95% reduction in FLOPs and 35%-80% reduction in memory while only 5%-51% reduction in FLOPs and 1%-17% reduction in memory are achieved by SNIP NP.

¹ We concretely define “resource” as Floating Point Operations per second (FLOPs) and memory required by one forward pass.

The main idea of RANP is to prune based on a *neuron importance* criterion analogous to the connection sensitivity in SNIP. Note that, pruning based on such a simple criterion as SNIP has the risk of pruning the whole layer(s) at extreme sparsity levels especially on large networks Lee et al. (2020). Even though an orthogonal initialization that ensures layer-wise dynamical isometry is sufficient to mitigate this issue for parameter pruning on 2D CNNs Lee et al. (2020), it is unclear if this could be directly applied to neuron pruning on 3D CNNs. To tackle this and improve pruning, we introduce a *resource aware reweighting scheme* that first balances the mean value of neuron importance in each layer and then reweights the neuron importance based on the resource consumption of each neuron. As evidenced by our experiments, such a reweighting scheme is crucial to obtain large reductions in memory and FLOPs while maintaining high accuracy.

We first evaluate our RANP on 3D semantic segmentation on a sparse point-cloud dataset, ShapeNet Yi et al. (2016), and a dense medical image dataset, BraTS’18 Menze et al. (2015); Bakas et al. (2017), with widely used 3D-UNets Cicek et al. (2016). We also evaluate RANP on video classification using UCF101 Soomro et al. (2012) with MobileNetV2 Sandler et al. (2018) and I3D Carreira and Zisserman (2017) and two-view stereo matching using Scene-Flow Mayer et al. (2016) with PSM Chang and Chen (2018) which has a high combination of 2D/3D convolution layers (roughly 3:1), see details in “3D CNNs” in Sec. 6.1. Our RANP-f significantly outperforms other neuron pruning methods in *resource efficiency* by yielding large reductions in computational resources (50%-95% FLOPs reduction and 35%-80% memory reduction) with comparable accuracy to the unpruned network (ref. Fig. 1).

We perform extensive experiments to demonstrate

- **Scalability** of RANP by pruning with a small input spatial size and training with a large one
- **Transferability** by pruning using ShapeNet and training on BraTS’18 and vice versa
- **Lightweight training** on a single GPU
- **Fast training** with increased batch size

2 Related Work

Previous works of network pruning mainly focus on 2D CNNs by parameter pruning Lee et al. (2019); Han et al. (2015); Guo et al. (2016); Dong et al. (2017); Chen et al. (2018) and neuron pruning Yu et al. (2017); He et al. (2017); Li et al. (2016); Yu et al. (2018); Huang and Wang (2018); He et al. (2018). While a majority of the pruning methods use the traditional prune-retrain scheme with a combined loss function of pruning criteria Han et al. (2015); He et al. (2017); Yu et al. (2018), some pruning at initialization methods is able

to reduce computational complexity in training [Lee et al. \(2019\)](#); [Zhang and Stadie \(2019\)](#); [Li et al. \(2019\)](#); [Yu and Huang \(2019\)](#); [Wang et al. \(2020\)](#). While very few are for 3D CNNs [Molchanov et al. \(2017\)](#); [Zhang et al. \(2019\)](#); [Chen et al. \(2020\)](#), none of them prune networks at initialization, and thus, none of them effectively reduce the training-time computational and memory requirements of 3D CNNs.

2D CNN pruning. *Parameter pruning* merely sparsifies filters for a high learning capability with small models. [Han et al. \(2015\)](#) used an iterative method of removing parameters with values below a threshold. [Lee et al. \(2019\)](#) recently proposed a single-shot method with connection sensitivity by magnitudes of parameter mask gradients to retain top- κ parameters. These filter-sparse methods, however, do not directly yield large memory reductions and speedup with decreased FLOPs.

By contrast, *neuron pruning*, also known as filter pruning or channel pruning, can effectively reduce computational resources. For instance, [Li et al. \(2016\)](#) used l_1 normalization to remove unimportant filters with connecting features. [He et al. \(2017\)](#) used a LASSO regression to prune network layers with reconstruction in the least square manner. [Yu et al. \(2017\)](#) proposed a group-wise 2D-filter pruning from each 3D-filter by a learning-based method and a knowledge distillation. Structure learning based MorphNet [Gordon et al. \(2018\)](#) and SSL [Wen et al. \(2016\)](#) aim at pruning activations with structure constraints or regularization. These approaches only reduce the test-time resource requirement while we focus on reducing the resource requirement of large 3D CNNs at training time.

3D CNN pruning. To improve the efficiency on 3D CNNs, some works like SSC [Graham et al. \(2018\)](#) and OctNet [Riegler et al. \(2017\)](#) use efficient data structures to reduce the memory requirement for sparse point-cloud data. However, these approaches are not useful for dense data, *e.g.*, MRI images and videos, and the resource requirement remains prohibitively large, tackling efficient network training.

Hence, it is desirable to develop an efficient pruning for 3D CNNs that can handle dense 3D data which is common in real applications. Only very few works are relevant to 3D CNN pruning. [Molchanov et al. \(2017\)](#) proposed a greedy criteria-based method to reduce resources via backpropagation with a small 3D CNN for hand gesture recognition. [Zhang et al. \(2019\)](#) used a regularization based pruning method by assigning regularization to weight groups with $4\times$ speedup in theory. Recently, [Chen et al. \(2020\)](#) converted 3D filters into frequency domain to eliminate redundancy in an iterative optimization for convergence. Being a parameter pruning method, this does not lead to large FLOPs and memory reductions, *e.g.*, merely a $2\times$ speedup compared to our $28\times$ (ref. Sec. 6.3). In summary, these methods embed pruning in the iterative net-

work optimization and require extensive resources, which is inefficient for 3D CNNs.

Pruning at Initialization. While few works used pruning at initialization, some achieved impressive success. SNIP [Lee et al. \(2019\)](#) is the first single-shot pruning method that presented a high possibility of pruning networks at initialization with minimal accuracy loss in training, followed by many recent works on single-shot pruning [Zhang and Stadie \(2019\)](#); [Li et al. \(2019\)](#); [Yu and Huang \(2019\)](#); [Wang et al. \(2020\)](#). But none are for 3D CNNs pruning.

In addition to being a parameter pruning approach, the benefits of SNIP was demonstrated only on small-scale datasets [Lee et al. \(2019\)](#), such as MNIST and CIFAR-10. Therefore, it is unclear that whether these benefits could be transposed to 3D CNNs applied to large-scale datasets. Our experiments indicate that, while SNIP itself is not capable of yielding large resource reduction on 3D CNNs, our RANP can greatly reduce the computational resources without causing network infeasibility. Furthermore, we show that RANP enjoys strong transferability among datasets and enables fast and lightweight training of large 3D volumetric data segmentation on a single GPU.

3 Preliminaries

We first briefly describe the main idea of SNIP [Lee et al. \(2019\)](#) which removes redundant parameters prior to training. Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^S$ with input \mathbf{x}_i and ground truth \mathbf{y}_i and the sparsity level κ , the optimization problem associated with SNIP can be written as

$$\begin{aligned} \min_{\mathbf{c}, \mathbf{w}} L(\mathbf{c} \odot \mathbf{w}; \mathcal{D}) &= \min_{\mathbf{c}, \mathbf{w}} \frac{1}{S} \sum_{i=1}^S \ell(\mathbf{c} \odot \mathbf{w}, (\mathbf{x}_i, \mathbf{y}_i)), \\ \text{s.t. } \mathbf{w} \in \mathbb{R}^m, \quad \mathbf{c} \in \{0, 1\}^m, \quad \|\mathbf{c}\|_0 \leq \kappa, \end{aligned} \quad (1)$$

where \mathbf{w} is denoted a m -dimensional vector of parameters, \mathbf{c} is the corresponding binary mask on the parameters, $\ell(\cdot)$ is the standard loss function (*e.g.*, cross-entropy loss), and $\|\cdot\|_0$ denotes l_0 norm. The mask $c_j \in \{0, 1\}$ for parameter w_j denotes that the parameter is retained in the compressed model if $c_j = 1$ and otherwise it is removed. In order to optimize the above problem, they first relax the binary constraint on the masks such that $\mathbf{c} \in [0, 1]^m$. Then an importance function for parameter w_j is calculated by the normalized magnitude of the loss gradient over mask c_j as

$$s_j = \frac{|g_j(\mathbf{w}; \mathcal{D})|}{\sum_{k=1}^m |g_k(\mathbf{w}; \mathcal{D})|}, \quad g_j(\mathbf{w}; \mathcal{D}) = \left. \frac{\partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D})}{\partial c_j} \right|_{\mathbf{c}=1}. \quad (2)$$

Then, only top- κ parameters are retained based on the parameter importance, called connection sensitivity in [Lee et al.](#)

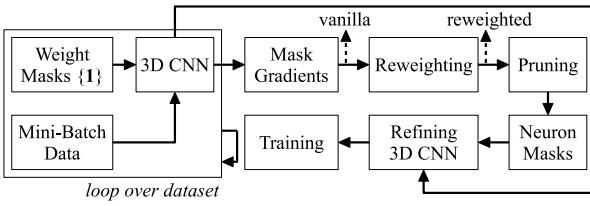


Fig. 2: Flowchart of RANP algorithm. The refining generates a new yet slim network for resource-efficient training.

(2019), s defined above. Upon pruning, the retained parameters are trained in the standard way. It is interesting to note that, even though having the mask \mathbf{c} is easier to explain the intuition, SNIP can be implemented without these additional variables by noting that $g_j(\mathbf{w}; \mathcal{D}) = (\partial L(\mathbf{w}; \mathcal{D}) / \partial w_j) w_j$ Lee et al. (2020). This method has shown remarkable results in achieving $> 95\%$ sparsity on 2D image classification tasks with minimal loss of accuracy. Such a parameter pruning method is important, however, it cannot lead to sufficient computation and memory reductions to train a deep 3D CNN on current off-the-shelf graphics hardware. In particular, the sparse weight matrices cannot efficiently reduce memory or FLOPs, and they require specialized sparse matrix implementations for speedup. In contrast, neuron pruning directly translates into practical gains of reducing both memory and FLOPs. This is crucial in 3D CNNs due to their substantially higher resource requirement compared to 2D CNNs.

4 Resource Aware NP at Initialization

To explain the proposed RANP, we first extend the SNIP idea to neuron pruning at initialization. Then we discuss a resource aware reweighting strategy to further reduce computational requirements of the pruned network. The flowchart of our RANP algorithm is shown in Fig. 2.

Before introducing our neuron importance, we first consider a fully-connected feed-forward neural network for the simplicity of notations. Consider weight matrices $\mathbf{W}^l \in \mathbb{R}^{N_l \times N_{l-1}}$, biases $\mathbf{b}^l \in \mathbb{R}^{N_l}$, pre-activations $\mathbf{h}^l \in \mathbb{R}^{N_l}$, and post-activations $\mathbf{x}^l \in \mathbb{R}^{N_l}$, for layer $l \in \mathcal{K} = \{1, \dots, K\}$. Now the feed-forward dynamics is

$$\mathbf{x}^l = \phi(\mathbf{h}^l), \quad \text{where } \mathbf{h}^l = \mathbf{W}^l \mathbf{x}^{l-1} + \mathbf{b}^l, \quad (3)$$

where the activation function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ has elementwise nonlinearity and the network input is denoted by \mathbf{x}^0 . Now we introduce binary masks on neurons (*i.e.*, post-activations). The feed-forward dynamics is then modified to include this masking operation as

$$\mathbf{x}^l = \mathbf{c}^l \odot \phi(\mathbf{h}^l), \quad \text{where } \mathbf{c}^l \in \{0, 1\}^{N_l}, \quad \forall l \in \mathcal{K}, \quad (4)$$

where neuron mask $c_u^l = 1$ indicates neuron x_u^l is retained and otherwise pruned. Here, neuron pruning can be written as the following optimization problem

$$\begin{aligned} \min_{\mathbf{w}} L(\mathbf{c}, \mathbf{w}; \mathcal{D}) &= \min_{\mathbf{c}, \mathbf{w}} \frac{1}{S} \sum_{i=1}^S \ell(\mathbf{c}, \mathbf{w}; (\mathbf{x}_i, \mathbf{y}_i)), \\ \text{s.t. } \mathbf{w} \in \mathbb{R}^m, \quad \mathbf{c} \in \{0, 1\}^n, \quad \|\mathbf{c}\|_0 \leq \kappa, \end{aligned} \quad (5)$$

where n is the total number of neurons and $\ell(\mathbf{c}, \cdot; \cdot)$ denotes a standard loss function of the feed-forward mapping with neuron masks \mathbf{c} defined in Eq. 4. This can be easily extended to convolutional and skip-concatenation operations.

As removing neurons could largely reduce memory and FLOPs compared to merely sparsifying parameters, the core of our approach is benefited by removing redundant neurons from the model. We use an influence function concept developed for parameters to establish neuron importance through the loss function, to locate redundant neurons.

4.1 Neuron Importance

Note that, neuron importance can be derived from the SNIP-based parameter importance discussed in Sec. 3. Another approach is to directly define neuron importance as the normalized magnitude of the neuron mask gradients analogous to parameter importance.

4.1.1 Neuron Importance with Parameter Mask Gradients

The first approach to calculate neuron importance depends on the magnitude of parameter mask gradients, denoted as Magnitude of Parameter Mask Gradients (MPMG). Thus, the importance of neuron x_u^l is

$$s_u^l = f(|g_{u1}^l|, \dots, |g_{uN_{l-1}}^l|), \quad (6)$$

where $g_{uv}^l = \partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D}) / \partial c_{uv}^l$ with c_{uv}^l as the mask of parameter w_{uv}^l . Refer to Eq. 2. Here, $f(\cdot) : \mathbb{R}^{N_{l-1}} \rightarrow \mathbb{R}$ is a function mapping a set of values to a scalar. We choose $f(\cdot) = \text{sum}(\cdot)$ with alternatives, *i.e.*, mean and max functions, in Table 1. Now, we set neuron masks as 1 for neurons with top- κ largest neuron importance.

4.1.2 Neuron Importance with Neuron Mask Gradients

Another approach is to directly compute mask gradients on neurons and treat their magnitudes as neuron importance, denoted as Magnitude of Neuron Mask Gradients (MNMG). The neuron importance of x_u^l is calculated by

$$s_u^l = \left| \frac{\partial L(\mathbf{c}, \mathbf{w}; \mathcal{D})}{\partial c_u^l} \right|_{\mathbf{c}=1}. \quad (7)$$

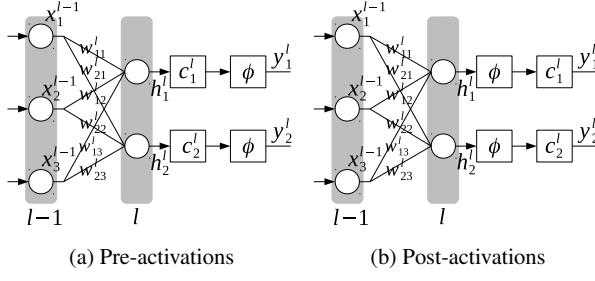


Fig. 3: Pre-activations and post-activations, where \mathbf{x} are layer inputs, \mathbf{w} are weights, \mathbf{c} are neuron masks, $\phi(\cdot)$ is an activation function, \mathbf{h} are hidden values, and \mathbf{y} are outputs.

Noting that a non-linear activation function $\phi(\cdot)$ in CNN including but not limited to ReLU can satisfy $\phi(ch) = c\phi(h)$, $\forall c \geq 0$. Given such a homogeneous function, the calculation of neuron importance with neuron masks can be derived from parameter mask gradients in the form of

$$\frac{\partial L(\mathbf{c}, \mathbf{w}; \mathcal{D})}{\partial c_u^l} \Big|_{\mathbf{c}=\mathbf{1}} = \sum_{v=1}^{N_{l-1}} \frac{\partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D})}{\partial c_{uv}^l} \Big|_{\mathbf{c}=\mathbf{1}}. \quad (8)$$

Details of the influence of such an activation function on neuron importance are provided in Sec. 4.1.3. These two approaches for neuron importance are in a similar form that while MPMG is by the sum of magnitudes, MNMG is by the magnitude of the sum of parameter mask gradients. It can be implemented directly from parameter gradients.

The neuron importance based on MPMG or MNMG approach can be used to remove redundant neurons. However, they could lead to an imbalance of sparsity levels of each layer in 3D network architectures. As shown in Table 2, the computational resources required by vanilla neuron pruning are much higher than those by other sparsity enforcing methods, e.g., random neuron pruning and layer-wise neuron pruning. We hypothesize that this is caused by the layer-wise imbalance of neuron importance which unilaterally emphasizes on some specific layer(s) and may lead to network infeasibility by pruning the whole layer(s). This behavior is also observed in Lee et al. (2020), and orthogonal initialization is thus recommended to solve the problem for 2D CNN pruning, which however cannot result in balanced neuron importance in our case, see results in Sec. 6.5.

In order to resolve this issue, we propose resource aware neuron pruning (RANP) with reweighted neuron importance, and the details are provided below.

4.1.3 Impacts of Activation Function

We first establish the relation between MPMG and MNMG for calculating neuron importance given a homogeneous activation function $\phi(\cdot)$ that includes but not limited to ReLU

used in the 3D CNNs. Then we analyze the impact of such an activation function on the calculation of neuron importance by deriving the mask gradients on post-activations and pre-activations illustrated in Figs. 3b and 3a respectively.

Proposition 1 *For a network activation function $\phi(w): \mathbb{R} \rightarrow \mathbb{R}$ being a homogeneous function of degree 1 satisfying $\phi(cw) = c\phi(w)$, $\forall c \geq 0$, the neuron mask gradient equals the sum of parameter mask gradients of this neuron.*

Proof: Given a neuron mask c_1 before the activation function $\phi(\cdot)$ in Fig. 3a and the output of the 1st neuron as y_1^l , we have

$$\begin{aligned} y_1^l &= \phi(c_1^l \odot h_1^l) \\ &= \phi(c_1^l \odot (x_1^{l-1} w_{11}^l + x_2^{l-1} w_{12}^l + x_3^{l-1} w_{13}^l)) \\ &= \phi(c_1^l x_1^{l-1} w_{11}^l + c_1^l x_2^{l-1} w_{12}^l + c_1^l x_3^{l-1} w_{13}^l). \end{aligned} \quad (9)$$

The gradient of loss L over the neuron mask c_1^l is

$$\begin{aligned} \frac{\partial L}{\partial c_1^l} &= \frac{\partial L}{\partial y_1^l} \frac{\partial y_1^l}{\partial c_1^l} \\ &= \frac{\partial L}{\partial y_1^l} (x_1^{l-1} w_{11}^l + x_2^{l-1} w_{12}^l + x_3^{l-1} w_{13}^l). \end{aligned} \quad (10)$$

Meanwhile, if setting masks on weights of this neuron directly, we can obtain

$$y_1^l = \phi(c_{11}^l x_1^{l-1} w_{11}^l + c_{12}^l x_2^{l-1} w_{12}^l + c_{13}^l x_3^{l-1} w_{13}^l), \quad (11)$$

then the gradient of weight mask, e.g., c_{11}^l , from loss is

$$\frac{\partial L}{\partial c_{11}^l} = \frac{\partial L}{\partial y_1^l} \frac{\partial y_1^l}{\partial c_{11}^l} = \frac{\partial L}{\partial y_1^l} x_1^{l-1} w_{11}^l. \quad (12)$$

Similarly,

$$\begin{aligned} \frac{\partial L}{\partial c_{12}^l} + \frac{\partial L}{\partial c_{13}^l} &+ \frac{\partial L}{\partial c_{11}^l} \\ &= \frac{\partial L}{\partial y_1^l} (x_1^{l-1} w_{11}^l + x_2^{l-1} w_{12}^l + x_3^{l-1} w_{13}^l). \end{aligned} \quad (13)$$

Clearly, Eq. 10 equals Eq. 13. Hence, the neuron mask gradients can be calculated by parameter mask gradients. To this end, the proof is done.

Furthermore, given such a homogeneous activation function in Prop. 1, the importance of a post-activation equals the importance of its pre-activation. In more detail, for post-activations in Fig. 3b, output y_1^l is

$$\begin{aligned} y_1^l &= c_1^l \odot \phi(h_1^l) \\ &= c_1^l \odot \phi(x_1^{l-1} w_{11}^l + x_2^{l-1} w_{12}^l + x_3^{l-1} w_{13}^l). \end{aligned} \quad (14)$$

Since the activation function satisfies $c\phi(w) = \phi(cw)$,

$$y_1^l = \phi(c_1^l x_1^{l-1} w_{11}^l + c_1^l x_2^{l-1} w_{12}^l + c_1^l x_3^{l-1} w_{13}^l). \quad (15)$$

The neuron importance determined by neuron mask c_1^l is

$$\begin{aligned} \frac{\partial L}{\partial c_1^l} &= \frac{\partial L}{\partial y_1^l} \frac{\partial y_1^l}{\partial c_1^l} \\ &= \frac{\partial L}{\partial y_1^l} (x_1^{l-1} w_{11}^l + x_2^{l-1} w_{12}^l + x_3^{l-1} w_{13}^l) . \end{aligned} \quad (16)$$

Clearly, Eq. 16 equals Eq. 10. Now, the importance of pre-activations and post-activations is the same given such a homogeneous activation function.

4.2 Weighting with Layerwise Neuron Importance

To tackle the imbalanced neuron importance issue above, we first weight the neuron importance across layers. Weighting neuron importance of x_u^l can be expressed as

$$\tilde{s}_u^l = \frac{\max_{k=1}^K \bar{s}^k}{\bar{s}^l} s_u^l, \text{ where } \bar{s}^k = \frac{1}{N_k} \sum_{u=1}^{N_k} s_u^k, \forall k \in \mathcal{K} . \quad (17)$$

Here, \bar{s}^l is the mean neuron importance of layer l and \tilde{s}_u^l is the updated neuron importance. This helps to achieve the same mean neuron importance in each layer, which largely avoids underestimating neuron importance of specific layer(s) to prevent from pruning the whole layer(s).

To further reduce the memory and FLOPs with minimal accuracy loss, we then reweight the neuron importance \tilde{s}_u^l by available resource, *i.e.*, memory or FLOPs. This reweighting counts on the addition of weighted neuron importance and the effect of the computational resource, denoted as RANP-[m|f], where “m” is for memory and “f” is for FLOPs.

4.3 Reweighting with Resource Constraints

Given input dimension of the l th layer (x_{in}, x_h, x_w, x_d) ², neuron dimension $(f_{out}, f_{in}, f_h, f_w, f_d)$, and output dimension (y_{in}, y_h, y_w, y_d) with $x_{in} = f_{in}$ and $f_{out} = y_{in}$, the resource importance of FLOPs or memory is defined by

$$\begin{aligned} \text{FLOPs: } \tau_l &= [(f_h f_w f_d + f_h f_w f_d - 1) f_{in} \\ &\quad + f_{in} - 1 + 1|_{\text{bias}}] y_{in} y_h y_w y_d \\ &= (2f_h f_w f_d f_{in} - 1 + 1|_{\text{bias}}) y_{in} y_h y_w y_d , \end{aligned} \quad (18a)$$

$$\text{Memory: } \tau_l = y_{in} y_h y_w y_d , \quad (18b)$$

where $(f_h f_w f_d)$ is the number of operations of multiplications of filter³ and layer input, $(f_h f_w f_d - 1)$ is for additions of values from the multiplications, (f_{in}) is for multiplications over all f_{in} filters, $(f_{in} - 1)$ is for additions of values from all these multiplications, $(1|_{\text{bias}})$ is for an addition

² The dimension order follows that of PyTorch.

³ Here, we refer a 3D filter with dimension (f_h, f_w, f_d) .

when the neuron has a bias, and $(y_{in} y_h y_w y_d)$ is for all elements of the layer output.

The reweighted neuron importance of neuron x_u^l by following weighted addition variant RANP-[m|f] is

$$\hat{s}_u^l = (1 + \lambda \text{softmax}(-\tau_l)) \tilde{s}_u^l = \left(1 + \lambda \frac{e^{-\tau_l}}{\sum_{k=1}^K e^{-\tau_k}} \right) \tilde{s}_u^l , \quad (19)$$

where coefficient $\lambda > 0$ helps to control the effect of resource on neuron importance. This effect represented by softmax constrains the values into a controllable range [0,1], making it easy to determine λ and function a high resource influence with a small resource occupation.

We demonstrate the effect of this reweighting strategy over vanilla pruning in Fig. 4. Take ShapeNet case as an example, vanilla neuron importance tends to have high values in the last few layers, making it highly possible to remove all neurons of such as the 7th and 8th layers. Weighting the importance in Fig. 4b makes the distribution of importance balanced with the same mean value in each layer. Furthermore, since some neurons have different numbers of input channels, each layer requires different FLOPs and memory. Considering the effect of computational resources on training, we embed them into neuron importance as weights.

In Fig. 4c, the last few layers require larger computational resources than the others, and thus their neurons share lower weights, see the tendency of mean values. Vividly, neuron ratio in Fig. 4d indicates a more balanced distribution by RANP-f than vanilla NP. For instance, very few neurons are retained in the 8th layer by vanilla NP, resulting in low accuracy and low maximum neuron sparsity. With reweighting by RANP-f, however, more neurons can be retained in this layer. Moreover, in Table 2, while weighted NP achieves high accuracy, its computational resource reductions are small. In contrast, RANP-f largely decreases the computational resources with a small accuracy loss.

Then, with reweighted neuron importance by Eq. 19 and \ddot{s}_κ as the κ th reweighted neuron importance in a descending order, the binary mask of neuron x_u^l can be obtained by

$$c_u^l = 1[\hat{s}_u^l - \ddot{s}_\kappa \geq 0] . \quad (20)$$

As mentioned in Sec. 2, our RANP is more effective in reducing memory and FLOPs than SNIP-based pruning which merely sparsifies parameters but needs high memory required by dense operations in training. RANP can easily remove neurons and all involved input channels at once, leading to huge reductions of input and output channels of the filter.

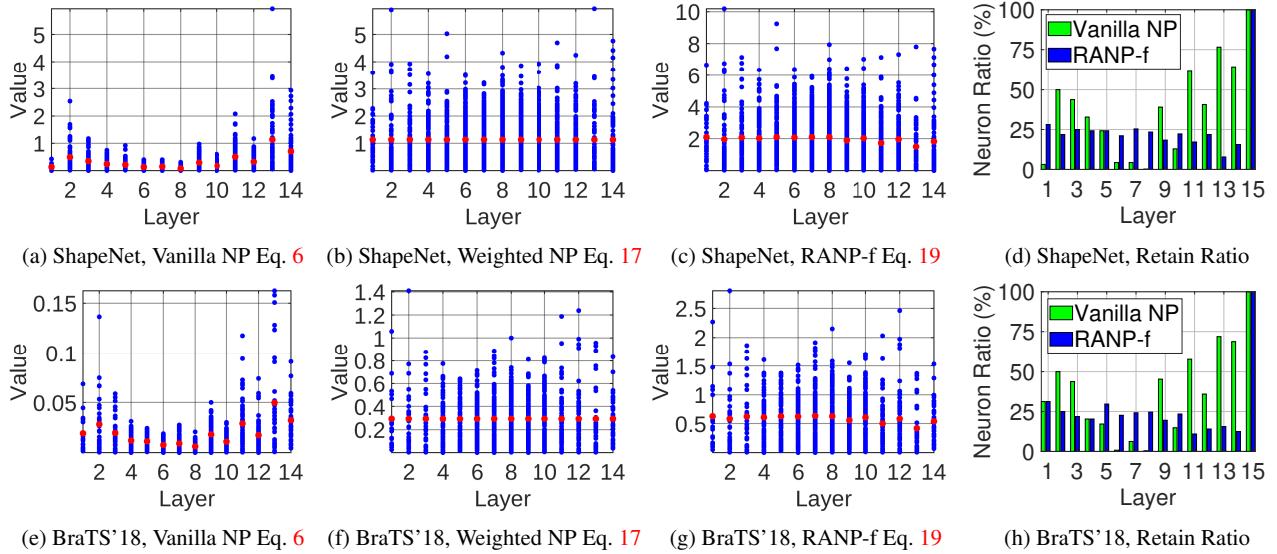


Fig. 4: Balanced neuron importance distribution by MPMG-sum on ShapeNet and BraTS'18. Neuron sparsity is 78.24% on ShapeNet and 78.17% on BraTS'18. Blue: neuron values; red: mean values.

5 Pseudocode of RANP Procedures

In Alg. 1, we provide the pseudocode of the pruning procedures of RANP. In Alg. 2, we used a simple half-space method to automatically search for the max neuron sparsity with network feasibility. Note that this searching cannot guarantee a small accuracy loss but merely to decide the maximum pruning capability.

6 Experiments

We evaluated RANP on 3D-UNets for 3D semantic segmentation and MobileNetV2 and I3D for video classification. Experiments are on Nvidia Tesla P100-SXM2-16GB GPUs in PyTorch. Our code is available at <https://github.com/zwxu064/RANP.git>.

6.1 Experimental Setup

3D Datasets. For 3D semantic segmentation, we used the large-scale 3D sparse point-cloud dataset, ShapeNet [Yi et al. \(2016\)](#), and dense biomedical MRI sequences, BraTS'18 [Menze et al. \(2015\); Bakas et al. \(2017\)](#).

ShapeNet consists of 50 object part classes, 14007 training samples, and 2874 testing samples. We split it into 6955 training samples and 7052 validation samples as [Graham et al. \(2018\)](#) to assign each point/voxel with a part class.

BraTS'18 includes 210 High Grade Glioma (HGG) and 75 Low Grade Glioma (LGG) cases. Each case has 4 MRI sequences, *i.e.*, T1, T1_CE, T2, and FLAIR. The task is to

detect and segment brain scan images into 3 categories: Enhancing Tumor (ET), Tumor Core (TC), and Whole Tumor (WT). The spatial size is $240 \times 240 \times 155$ in each dimension. We used the splitting strategy of cross-validation in [Kao et al. \(2018\)](#) with 228 cases for training and 57 cases for validation.

For video classification, we used *UCF101* [Soomro et al. \(2012\)](#) with 101 action categories and 13320 videos. 2D spatial dimension from images and temporal dimension from frames are cast as dense 3D inputs. Among the 3 official train/test splits, we used split-1 which has 9537 videos for training and 3783 videos for validation.

For two-view stereo matching, we used *SceneFlow* [Mayer et al. \(2016\)](#) containing 3 synthetic videos flyingthings3D, driving, and Monkaa for 192-disparity estimation. In total, 35,454 image pairs were used for training and 4,370 image pairs for validation. The image pairs have a high resolution of 540×960 pixels in RGB format.

3D CNNs. For 3D semantic segmentation on ShapeNet (sparse data) and BraTS'18 (dense data), we used the standard 15-layer 3D-UNet [Cicek et al. \(2016\)](#) including 4 encoders, each consists of two “3D convolution + 3D batch normalization + ReLU”, a “3D max pooling”, four decoders, and a confidence module by softmax. It has 14 convolution layers with 3^3 kernels and 1 layer with 1^3 kernel.

For video classification, we used the popular *MobileNetV2* [Sandler et al. \(2018\); Kopuklu et al. \(2019\)](#) and *I3D* (with inception as backbone) [Carreira and Zisserman \(2017\)](#) on UCF101. MobileNetV2 has a linear layer and 52 convolution layers while 18 of them are 3^3 kernels and the rest are

Algorithm 1: Pruning Procedures of RANP-[f|m].

Input: Dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^S$ with B samples per batch, neuron sparsity κ , resource importance $\{\tau_l\}$, coefficient $\lambda > 0$, and parameter masks $\mathbf{c} = \{c_{uv}^l\}$, where layer $l \in \mathcal{K} = \{1, \dots, K\}$, and neuron $u \in \mathcal{N}_l = \{1, \dots, N_l\}$.

Output: Binary neuron masks $\hat{\mathbf{c}} = \{\hat{c}_u^l\}$.

```

1 for batch  $t \in \{1, \dots, \lfloor S/B \rfloor\}$  do                                 $\triangleright$  mini-batch
2    $\mathcal{D}^t \leftarrow \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=(t-1)B+1}^{tB}$                           $\triangleright$  parameter mask gradient, Eq. 2
3    $g_{uv}^l \leftarrow \partial L(\mathbf{c} \odot \mathbf{w}; \mathcal{D}^t) / \partial c_{uv}^l$             $\triangleright$  parameter mask importance, Eq. 6
4    $g_{uv}^l \leftarrow |g_{uv}^l|$ , for MPMG                                      $\triangleright$  gradient accumulation
5    $\nabla c_{uv}^l \leftarrow g_{uv}^l$ ,  $\forall u \in \mathcal{N}_l, \forall v \in \mathcal{N}_{l-1}$            $\triangleright$  average on mini-batch
6    $\nabla c_{uv}^l \leftarrow \nabla c_{uv}^l / \lfloor S/B \rfloor$ ,  $\forall u \in \mathcal{N}_l, \forall v \in \mathcal{N}_{l-1}, \forall l \in \mathcal{K}$      $\triangleright$  vanilla neuron importance, Eq. 8
7    $s_u^l \leftarrow |\sum_{v=1}^{N_{l-1}} \nabla c_{uv}^l|$ ,  $\forall u \in \mathcal{N}_l, \forall l \in \mathcal{K}$             $\triangleright$  mean neuron importance, Eq. 17
8    $\bar{s}^l \leftarrow \sum_{u \in \mathcal{N}_l} s_u^l / N_l$ ,  $\forall l \in \mathcal{K}$                             $\triangleright$  weighting, Eq. 17
9    $\tilde{s}_u^l \leftarrow (\max_{j \in \mathcal{K}} \bar{s}^j / \bar{s}^l) s_u^l$ ,  $\forall u \in \mathcal{N}_l, \forall l \in \mathcal{K}$        $\triangleright$  reweighting, Eq. 19
10   $\hat{s}_u^l \leftarrow (1 + \lambda e^{-\tau_l} / \sum_{j \in \mathcal{K}} e^{-\tau_j}) \tilde{s}_u^l$ ,  $\forall u \in \mathcal{N}_l, \forall l \in \mathcal{K}$      $\triangleright$  sorting in descending
11   $\{\hat{s}_u\} \leftarrow \text{SortDescending}(\{\hat{s}_u^l\})$ ,  $\forall u \in \mathcal{N}_l, \forall l \in \mathcal{K}$             $\triangleright$  sorting in descending
12   $\hat{c}_u^l \leftarrow 1[\hat{s}_u^l - \hat{s}_\kappa \geq 0]$ ,  $\forall u \in \mathcal{N}_l, \forall l \in \mathcal{K}$             $\triangleright$  binary neuron mask, Eq. 20

```

Algorithm 2: Auto-Search for Max Neuron Sparsity .

Input: Dataset \mathcal{D} , layerwise resource usage τ w.r.t. FLOPs or memory, coefficient $\lambda > 0$, lower and upper sparsity κ_{min} and κ_{max} , threshold $\delta = 1e-4$. “feasible network” means not all neurons are removed in each layer.

Output: Max neuron sparsity κ^* .

```

1 Initialize  $\kappa_{min} \leftarrow 0, \kappa_{max} \leftarrow 1$ 
2 while  $(\kappa_{max} - \kappa_{min} > \delta)$  do
3    $\kappa = 0.5(\kappa_{min} + \kappa_{max})$ 
4    $y = \text{NeuronPruning}(\mathcal{D}, \tau, \lambda, \kappa)$                                 $\triangleright$  Alg. 1
5   if  $y == 0$  (feasible network) then
6      $\kappa_{min} \leftarrow \kappa$ 
7   else
8      $\kappa_{max} \leftarrow \kappa$ 
9    $\kappa^* = \kappa$ 

```

¹³. I3D has a linear layer and 57 convolution layers, 19 of which are 3^3 kernels, 1 is 7^3 , and the rest are 1^3 .

For two-view stereo matching, we used the widespread PSM Chang and Chen (2018) network, which has multi-scale feature extraction with 2D convolution layers for left-right image feature cost volumes, followed by 3 encoder-decoder modules with 3D convolution layers. PSM highly combines 2D/3D convolution layers with 61 2D layers (including linear layer) and 28 3D layers where 4,352 neurons are in the 2D layers and 1,283 neurons in the 3D layers.

In PSM, massive cross-layer additions are adopted to preserve high-level features. Three losses from the 3D encoder-decoder modules are weighted for training. Specifically, for the massive additions across multiple layer outputs in PSM, we used the maximum retained channels as the out_channels for all associated layers. For instance, given 3 layers having 256 out_channels (also denoted as neurons) for addition, they retain 100, 50, 120 out_channels respectively after pruning. To maintain the addition, they are required to have

the same number of out_channels. We thus selected the maximum number 120 by $\max\{100, 50, 120\}$ as out_channels of these 3 layers and in_channels of the following layer in the refinement for a lightweight network.

Hyper-Parameters in Learning. For *ShapeNet*, we set learning rate as 0.1 with an exponential decay rate $\gamma = 0.04$ by 100 epochs; batch size is 12 on 2 GPUs; spatial size for pruning and training is 64^3 while the spatial size for training is 128^3 in Sec. 6.8; optimizer is SGD-Nesterov Sutskever et al. (2013) with weight decay 0.0001 and momentum 0.9.

For *BraTS’18*, learning rate is 0.001, decayed by 0.1 at 150th epoch with 200 epochs; optimizer is Adam Kingma and Ba (2015) with weight decay 0.0001 and AMSGrad Reddi et al. (2018); batch size is 2 on 2 GPUs; spatial size for pruning is 96^3 and 128^3 for training.

For *UCF101*, we used similar setup from Kopuklu et al. (2019) with learning rate 0.1, decayed by 0.1 at {40, 55, 60, 70}th epoch; optimizer by SGD with weight decay 0.001; batch size 8 on one GPU. Spatial size for pruning and training is 112^2 for MobileNetV2 and 224^2 for I3D; 16 frames are used for the temporal size. Note that in Kopuklu et al. (2019) networks for UCF101 had higher performance since they were pretrained on Kinetics600, while we directly trained on UCF101. A feasible train-from-scratch reference could be Soomro et al. (2012).

For *SceneFlow*, we trained PSM with 15 epochs, learning rate 0.001, 192 disparities, Adam optimizer, and 12 batch size on 3 GPUs. Crop size for training is 256×512 (height \times width) from the original size 540×960 .

For Eq. 19, we empirically set the coefficient λ as 11 for ShapeNet, 15 for BraTS’18, 80 for UCF101, and 100 for SceneFlow. Glorot initialization Glorot and Bengio (2010) was used for weight initialization. Note that we used orthogonal initialization Saxe et al. (2014) to handle imbalanced layer-wise neuron importance distribution Lee et al. (2020) but obtained lower maximum neuron sparsity.

Loss Function and Metrics. Due to the page limitation, we provide loss functions and metrics used in our experiments. Standard cross-entropy function was used as the loss function for *ShapeNet* and *UCF101*. For *Brats'18*, the weighted function in Kao et al. (2018) is

$$L = L_{ce} + \alpha L_{dice} = L_{ce} + \alpha \frac{1}{C} \sum_{i=1}^C \frac{2|\mathbf{P}_i \cap \mathbf{G}_i|}{|\mathbf{P}| + |\mathbf{G}|}, \quad (21)$$

where $\alpha = 0.25$ is an empiric weight for dice loss, \mathbf{P} is prediction, \mathbf{G} is ground truth, and C is the number of classes. For *SceneFlow*, the standard smooth L-1 loss was used in each of the 3 encoder-decoder modules in PSM Chang and Chen (2018) with weights 0.5, 0.7, and 1.0 in training phase.

$$L(d^p, d^g) = \frac{1}{N} \sum_{i=1}^N \text{SmoothL-1}(d_i^p - d_i^g), \quad (22)$$

with

$$\text{SmoothL-1}(x) = \begin{cases} 0.5x^2 & \text{if } |x| < 1 \\ |x| - 0.5 & \text{otherwise} \end{cases}, \quad (23)$$

where d_i^p and d_i^g are predicted disparity and ground truth disparity at pixel i respectively, and N is the number of image pixels without occlusion.

Meanwhile, *ShapeNet* accuracy was measured by mean IoU over each part of object category Yi et al. (2017) while IoU by $|\mathbf{P} \cap \mathbf{G}| / |\mathbf{P} \cup \mathbf{G}|$ was adopted for *Brats'18*. For *UCF101* video classification, top-1 and top-5 recall rates were used. For *SceneFlow*, non-occluded ground truth disparity maps were used for evaluation. End-Point-Error (EPE) was calculated by the absolute pixel-wise disparity difference between the ground truth and the prediction without occlusion areas. Accuracy was measured by the absolute pixel-wise disparity difference under a given pixel threshold, 1 pixel for Acc-1 and 3 pixels for Acc-3 in our case.

6.2 Selection for Vanilla Neuron Pruning

Here, we selected sum operation in MPMG and MNMG for vanilla neuron pruning considering the trade-off between computational resources and accuracy. To give a comprehensive study, we demonstrate detailed results of mean, max, and sum operations of MPMG and MNMG in Table 1. We relaxed the sum operation in Eq. 8 to mean and max.

In Table 1, we aim at obtaining the maximum neuron sparsity due to the target of reducing the computational resources at an extreme sparsity level with minimal accuracy loss. Vividly, for *ShapeNet*, MPMG-sum achieves the largest maximum neuron sparsity 78.24% among all with merely $\sim 0.53\%$ accuracy loss. Differently, for *Brats'18*, MNMG-sum has the largest maximum neuron sparsity 81.32%; however, the accuracy loss can reach up to $\sim 8.48\%$. In contrast,

while MPMG-sum has the second-largest maximum neuron sparsity 78.17%, the accuracy loss is much smaller than MNMG-sum. For *UCF101*, it is surprising that many manners have low accuracy. As we analyse the reason in the footnote in Table 1, with the extreme neuron sparsity, some layers of the pruned networks have only 1 neuron retained, losing sufficient features for learning, and thus, leading to low accuracy. For *SceneFlow*, MPMG-sum achieves the highest accuracy and EPE among all pruning methods. The resource reductions, i.e., 25.92% FLOPs and 7.68% memory, however, are much smaller than MNMG-sum, i.e., 41.34% FLOPs and 21.67% memory, due to the small maximum neuron sparsity 20.92%. This is possibly caused by the massive output addition and concatenations across multiple layers and a high combination of 2D and 3D convolution layers.

Hence, considering the comprehensive performance of reducing resources and maintaining the accuracy, MPMG-sum is selected as vanilla NP for general 3D CNNs (without massive cross-additions, cross-concatenations, and a high combination of 2D CNNs) and MNMG-sum is advised to be compared with MPMG-sum for the trade-off between resource reductions and accuracy. Note that any neuron sparsity greater than the maximum neuron sparsity will make the pruned network infeasible by pruning the whole layer(s).

6.3 Evaluation of RANP on Pruning Capability

Random NP retains κ neurons with neuron indices randomly shuffled. Layer-wise NP retains neurons using the same retain rate as κ in each layer. For SNIP-based parameter pruning, the parameter masks are post processed by removing redundant parameters and then making sparse filters dense, which is denoted as SNIP NP. For a fair comparison with SNIP NP, we used the maximum parameter sparsity 98.98% for *ShapeNet*, 98.88% for *Brats'18*, 86.26% for *MobileNetV2*, and 81.09% for *I3D*.

ShapeNet. Compared with random NP and layer-wise NP in Table 2, the maximum reduced resources by vanilla NP are much less due to the imbalanced layer-wise distribution of neuron importance. Weighted neuron importance by Eq. 17, however, further reduces 18.3% FLOPs and 29.6% memory with 0.14% accuracy loss.

Reweighting by RANP-f and RANP-m further reduces FLOPs and memory on the basis of weighted NP. Here, RANP-f can reduce 96.8% FLOPs, 95.3% parameters, and 73.7% memory over the unpruned networks. Furthermore, with a similar resource in Table 3, RANP achieves $\sim 0.5\%$ increase in accuracy. Note that a too-large λ can additionally reduce the resources but at the cost of accuracy. We illustrate 4 point-cloud semantic segmentation cases with spatial size 64^3 in Fig. 5.

Table 1: Vanilla neuron pruning method. **Main resource consumption** (GFLOPs and memory) are considered but not parameters whose resource consumption is much smaller than memory. Among the neuron pruning methods, we mark bold **the best** and underline **the second best**. “ \uparrow ” and “ \downarrow ” in “Metrics” denote the larger and the smaller the better respectively. “EPE”: End-Point Error. Overall, we selected MPMG-sum as vanilla NP for ShapeNet, BraTS’18, and UCF101 and MNMG-sum as vanilla NP for SceneFlow; the corresponding neuron sparsity for large resource reductions with small accuracy loss.

Dataset	Model	Manner	Sparsity (%)	Param (MB)	GFLOPs	Memory (MB)	Metrics (% except pixel for EPE)
							mIoU \uparrow
ShapeNet	3D-UNet	Full	0	62.26	237.85	997.00	83.79 \pm 0.21
		MPMG-mean	68.10	5.08	110.14	819.97	83.33 \pm 0.18
		MPMG-max	70.24	4.54	107.38	809.88	83.79\pm0.10
		MPMG-sum	78.24	2.54	55.69	557.32	83.26 \pm 0.14
		MNMG-mean	63.03	4.23	112.95	834.98	83.46 \pm 0.13
		MNMG-max	73.93	3.67	103.57	796.44	83.51 \pm 0.08
		MNMG-sum	66.93	4.29	<u>100.34</u>	<u>783.14</u>	83.65 \pm 0.02
BraTS’18	3D-UNet	Full	0	15.57	478.13	3628.00	ET \uparrow TC \uparrow WT \uparrow
		MPMG-mean	65.64	1.48	226.86	3038.27	<u>73.51\pm0.82</u> <u>73.28\pm1.14</u> <u>87.15\pm0.43</u>
		MPMG-max	75.78	0.83	189.43	2812.53	73.67\pm0.98 72.73 \pm 1.70 86.44 \pm 0.71
		MPMG-sum	78.17	0.55	<u>104.50</u>	<u>1936.44</u>	71.94 \pm 1.68 69.39 \pm 2.29 84.68 \pm 0.78
		MNMG-mean	63.85	1.08	176.76	2790.64	73.35 \pm 0.70 73.38\pm0.94 87.21\pm0.38
		MNMG-max	80.05	0.59	169.99	2676.05	72.52 \pm 1.91 72.40 \pm 1.74 84.63 \pm 0.60
		MNMG-sum	81.32	0.35	73.50	1933.20	64.48 \pm 1.10 68.47 \pm 1.59 80.71 \pm 1.07
UCF101	MobileNetV2	Full	0	9.47	0.58	157.47	Top-1 \uparrow Top-5 \uparrow
		MPMG-mean	26.31	4.39	0.55	156.00	2.98 \pm 0.14 14.04 \pm 0.14
		MPMG-max	29.48	3.96	0.54	155.38	3.49 \pm 0.12 13.64 \pm 0.10
		MPMG-sum	33.15	6.35	0.55	155.17	46.32\pm0.79 75.42\pm0.60
		MNMG-mean	38.91	2.79	<u>0.50</u>	<u>147.69</u>	<u>29.13\pm0.92</u> <u>62.93\pm1.37</u>
		MNMG-max	50.33	2.59	0.53	153.45	2.84 \pm 0.06 13.40 \pm 0.23
		MNMG-sum	39.89	4.66	0.43	120.01	1.03 \pm 0.00 5.76 \pm 0.00
	I3D	Full	0	47.27	27.88	201.28	51.58 \pm 1.86 77.35 \pm 0.63
		MPMG-mean	16.47	31.57	26.50	196.51	<u>51.88\pm2.00</u> 77.98 \pm 1.46
		MPMG-max	19.83	30.06	26.31	195.62	52.44\pm1.25 78.08\pm1.27
		MPMG-sum	25.32	29.93	25.76	192.42	51.57 \pm 1.46 <u>78.07\pm1.34</u>
		MNMG-mean	35.36	16.69	15.37	124.85	49.26 \pm 0.96 75.70 \pm 1.49
		MNMG-max	40.27	17.86	23.73	184.77	44.90 \pm 1.19 74.43 \pm 1.26
		MNMG-sum	32.87	20.00	<u>16.03</u>	<u>125.17</u>	46.90 \pm 1.26 74.02 \pm 1.25
SceneFlow	PSM	Full	0	19.93	771.02	7217.37	Acc-1 \uparrow Acc-3 \uparrow EPE \downarrow
		MPMG-mean	21.50	13.50	575.13	6689.85	85.97 \pm 0.70 94.21 \pm 0.33 1.42 \pm 0.07
		MPMG-max	27.25	11.61	540.05	6467.20	85.32 \pm 0.77 93.71 \pm 0.43 1.53 \pm 0.11
		MPMG-sum	20.92	14.03	571.21	6663.24	85.75\pm0.48 94.17\pm0.25 1.43\pm0.06
		MNMG-mean	44.33	6.00	449.20	<u>5744.31</u>	85.04 \pm 0.50 93.07 \pm 0.23 1.54 \pm 0.06
		MNMG-max	50.07	6.85	473.74	5872.91	85.11 \pm 0.71 93.69 \pm 0.36 1.53 \pm 0.09
		MNMG-sum	46.20	6.52	452.25	5653.44	85.24 \pm 0.77 93.83 \pm 0.41 1.51 \pm 0.09

BraTS’18. In Table 2, RANP-f achieves 96.5% FLOPs, 95.1% parameters, and 80% memory reductions. It further reduces 18.3% FLOPs and 33.3% memory over vanilla NP while increasing -1.21% ET, 5.11% TC, and 0.77% WT. With a similar resource in Table 3, RANP achieves higher accuracy than random NP and layer-wise NP.

Additionally, Chen et al. (2020) achieved 2 \times speedup on BraTS’18 with 3D-UNet. In comparison, our RANP-f has roughly 28 \times speedup, which is theoretically evidenced by the reduced FLOPs from 478.13G to 16.97G in Table 2. We illustrated 2 MRI cases in Fig. 6.

UCF101. In Table 2, for MobileNetV2, RANP-f reduces 55.2% FLOPs, 49% parameters, and 44.1% memory with around 1% accuracy loss. Meanwhile, for I3D, it reduces

49.9% FLOPs, 43.5% parameters, and 35.3% memory with around 2% accuracy increase. The RANP-based methods can reduce much more resources than other methods. We illustrated 2 video categories in Fig. 7.

SceneFlow. In Table 2, the proposed RANP-f reduces 60.2% parameters, 53.0% FLOPs, and 33.3% memory over the unpruned network with 0.74%-1.46% accuracy loss and 0.16 pixel EPE loss. Specifically, the reweighting scheme in RANP-f boosts the pruning ability of vanilla NP by roughly 10% reductions in FLOPs and memory. Our RANP-f achieves the highest resource reductions than SNIP NP, random NP, and layerwise NP with relatively small accuracy loss. We illustrated 4 examples of SceneFlow validation set in Fig. 8.

Table 2: Evaluation of neuron pruning capability. All models are **trained from scratch** for 100 epochs on ShapeNet and UCF101, 200 on BraTS’18, and 15 on SceneFlow. Metrics are calculated by the last 5 epochs. “sparsity” is the maximum parameter sparsity for SNIP NP and the maximum neuron sparsity for vanilla NP for others. Among the neuron pruning methods, we mark bold **the best** and underline **the second best**. “ \downarrow ” in resources denotes reduction in %; “ \uparrow ” and “ \downarrow ” in “Metrics” denote the larger and the smaller the better respectively. Overall, our RANP-f performs best with large reductions of **main resource consumption** (GFLOPs and memory) with negligible accuracy loss.

Dataset	Model	Manner	Sparsity (%)	Param (MB)	GFLOPs	Memory (MB)	Metrics (% except pixel for EPE)
ShapeNet	3D-UNet	ours	Full	0	62.26	237.85	997.00
			SNIP NP	98.98	5.31 (91.5 \downarrow)	126.22 (46.9 \downarrow)	833.20 (16.4 \downarrow)
			Random NP		3.05 (95.1 \downarrow)	10.36 (95.6 \downarrow)	267.95 (73.1 \downarrow)
			Layer-wise NP		2.99 (95.2 \downarrow)	11.63 (95.1 \downarrow)	296.22 (70.3 \downarrow)
			Vanilla NP	78.24	2.54 (95.9 \downarrow)	55.69 (76.6 \downarrow)	557.32 (44.1 \downarrow)
			Weighted NP		2.97 (95.2 \downarrow)	12.06 (94.9 \downarrow)	301.56 (69.8 \downarrow)
			RANP-m		3.39 (94.6 \downarrow)	6.68 (97.2\downarrow)	214.95 (78.4\downarrow)
			RANP-f		2.94 (95.3 \downarrow)	7.54 (96.8 \downarrow)	262.66 (73.7 \downarrow)
BraTS’18	3D-UNet	ours	Full	0	15.57	478.13	3628.00
			SNIP NP	98.88	1.09 (93.0 \downarrow)	233.11 (51.2 \downarrow)	2999.64 (17.3 \downarrow)
			Random NP		0.75 (95.2 \downarrow)	22.59 (95.3 \downarrow)	817.59 (77.5 \downarrow)
			Layer-wise NP		0.75 (95.2 \downarrow)	24.09 (95.0 \downarrow)	836.88 (77.0 \downarrow)
			Vanilla NP	78.17	0.55 (96.5 \downarrow)	104.50 (78.1 \downarrow)	1936.44 (46.6 \downarrow)
			Weighted NP		0.79 (95.0 \downarrow)	22.40 (95.3 \downarrow)	860.64 (76.3 \downarrow)
			RANP-m		0.87 (94.4 \downarrow)	13.47 (97.2\downarrow)	506.97 (86.0\downarrow)
			RANP-f		0.76 (95.1 \downarrow)	16.97 (96.5 \downarrow)	729.11 (80.0 \downarrow)
UCF101	MobileNetV2	ours	Full	0	9.47	0.58	157.47
			SNIP NP	86.26	3.67 (61.3 \downarrow)	0.54 (6.9 \downarrow)	155.35 (1.3 \downarrow)
			Random NP		4.58 (51.6 \downarrow)	0.34 (41.4 \downarrow)	106.68 (32.3 \downarrow)
			Layer-wise NP		4.56 (51.8 \downarrow)	0.33 (43.1 \downarrow)	106.92 (32.1 \downarrow)
			Vanilla NP	33.15	6.35 (32.9 \downarrow)	0.55 (5.2 \downarrow)	155.17 (1.5 \downarrow)
			Weighted NP		4.82 (49.1 \downarrow)	0.30 (48.3 \downarrow)	100.33 (36.3 \downarrow)
			RANP-m		4.87 (48.6 \downarrow)	<u>0.27 (53.4\downarrow)</u>	84.51 (46.3\downarrow)
			RANP-f		4.83 (49.0 \downarrow)	0.26 (55.2\downarrow)	88.01 (44.1 \downarrow)
SceneFlow	I3D	ours	Full	0	47.27	27.88	201.28
			SNIP NP	81.09	30.06 (36.4 \downarrow)	26.31 (5.6 \downarrow)	195.62 (2.8 \downarrow)
			Random NP		26.36 (44.2 \downarrow)	16.45 (41.0 \downarrow)	145.07 (27.9 \downarrow)
			Layer-wise NP		26.67 (43.6 \downarrow)	16.93 (39.3 \downarrow)	150.95 (25.0 \downarrow)
			Vanilla NP	25.32	29.93 (36.7 \downarrow)	25.76 (7.6 \downarrow)	192.42 (4.4 \downarrow)
			Weighted NP		26.57 (43.8 \downarrow)	15.56 (44.2 \downarrow)	142.57 (29.2 \downarrow)
			RANP-m		26.75 (43.4 \downarrow)	<u>14.08 (49.5\downarrow)</u>	<u>130.44 (35.2\downarrow)</u>
			RANP-f		26.69 (43.5 \downarrow)	13.98 (49.9\downarrow)	130.22 (35.3\downarrow)
SceneFlow	PSM	ours	Full	0	19.93	771.02	7217.37
			SNIP NP	90.66	11.09 (44.4 \downarrow)	512.97 (33.5 \downarrow)	6295.07 (12.8 \downarrow)
			Random NP		8.13 (59.2 \downarrow)	371.48 (51.8 \downarrow)	5014.83 (30.5 \downarrow)
			Layerwise NP		7.71 (61.3 \downarrow)	379.05 (50.8 \downarrow)	5009.78 (30.6 \downarrow)
			Vanilla NP	46.20	6.52 (67.3 \downarrow)	452.25 (41.3 \downarrow)	5653.44 (21.7 \downarrow)
			Weighted NP		7.90 (60.4 \downarrow)	386.66 (49.9 \downarrow)	5002.16 (30.7 \downarrow)
			RANP-m		7.96 (60.1 \downarrow)	353.78 (54.1\downarrow)	4693.55 (35.0\downarrow)
			RANP-f		7.93 (60.2 \downarrow)	362.77 (53.0 \downarrow)	4813.24 (33.3 \downarrow)
Acc-1 \uparrow Acc-2 \uparrow EPE \downarrow							

Table 3: In addition to Table 2, with similar GFLOPs or memory on 3D-UNets, our RANP-f achieves the highest accuracy.

Manner	ShapeNet					BraTS’18						
	Sparsity	Param	GFLOPs	Mem	mIoU	Sparsity	Param	GFLOPs	Mem	ET	TC	WT
Random NP	81.01	2.27	~7.54	253.12	82.66 \pm 0.23	81.08	0.56	~16.97	685.77	61.09 \pm 1.87	68.94 \pm 2.44	78.89 \pm 2.47
Layer-wise NP	82.82	1.84	~7.54	255.67	82.82 \pm 0.26	83.50	0.46	~16.97	700.64	70.50 \pm 0.63	74.27 \pm 0.95	83.63 \pm 0.92
Random NP	78.83	2.87	9.57	~262.66	82.86 \pm 0.45	80.90	0.57	17.95	~729.11	68.45 \pm 1.11	70.67 \pm 1.21	75.02 \pm 0.79
Layer-wise NP	82.81	1.94	8.14	~262.66	82.52 \pm 0.13	82.45	0.51	17.31	~729.11	70.45 \pm 1.03	69.27 \pm 1.95	82.42 \pm 0.68
RANP-f(ours)	78.24	2.94	7.54	262.66	83.07\pm0.22	78.17	0.76	16.97	729.11	70.73\pm0.66	74.50\pm1.05	85.45\pm1.06

Moreover, in Table 4, we compare the resource consumption of 2D and 3D CNNs in PSM which has a high combina-

tion of 2D/3D layers mentioned in “3D CNNs” in Sec. 6.1. Although in 2D CNNs, our RANP-f reduces less FLOPs

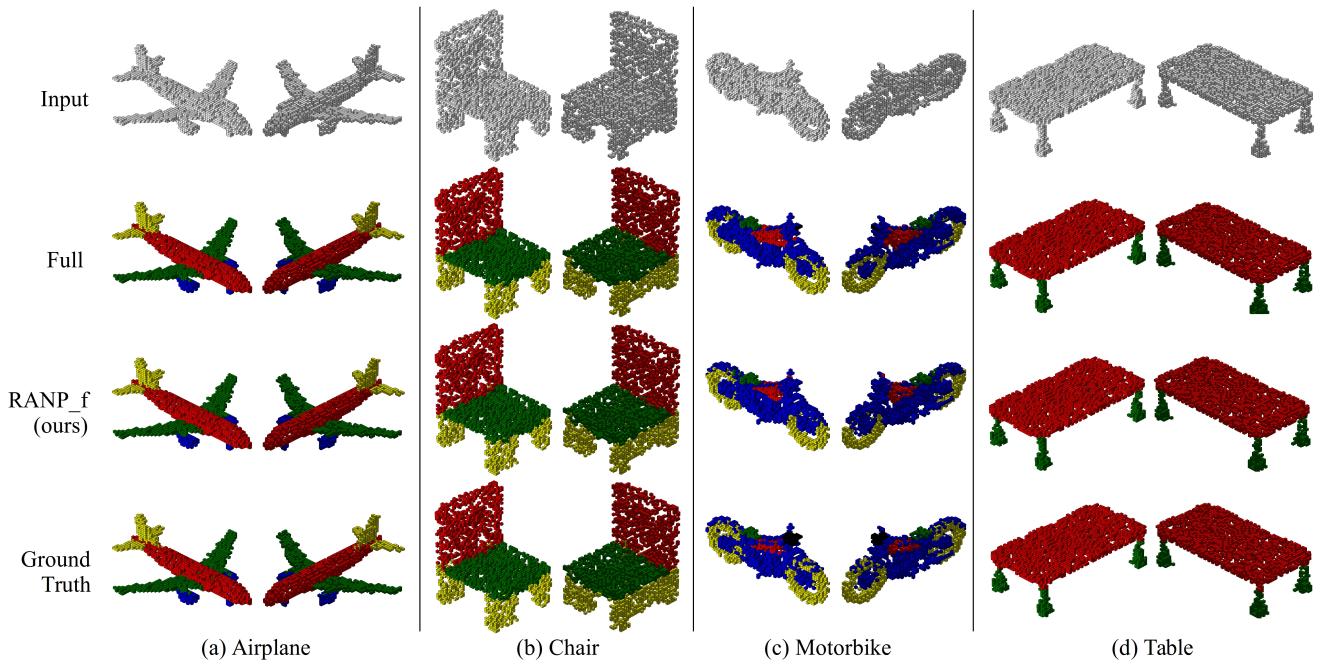


Fig. 5: Examples of ShapeNet for 3D semantic segmentation. Spatial size is 64^3 for the volumetric representation of sparse point clouds. We illustrate two views for each model, with (elevation angle, azimuth angle) as $(30^\circ, -45^\circ)$ and $(30^\circ, 45^\circ)$.

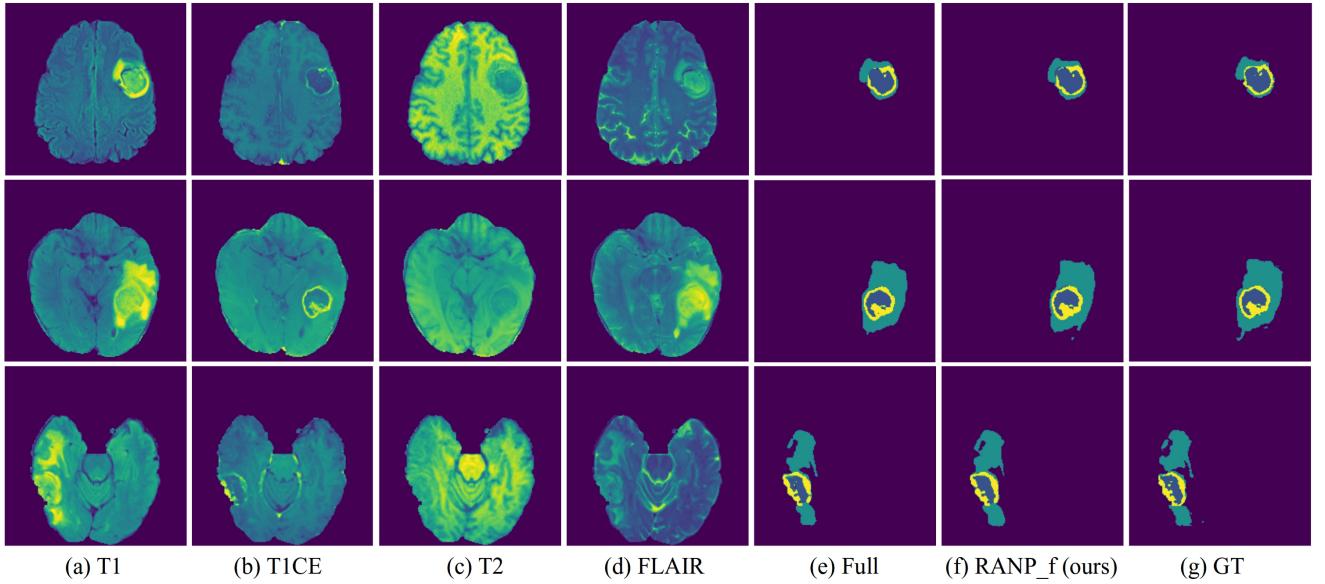


Fig. 6: Examples of BraTS'18 for medical image segmentation. Spatial size is 192^3 for the MRI image sequences. Slices with distinguishing segmentation are illustrated. (a)-(d) are input slices, (e)-(g) are segmentation with semantics: whole tumor, tumor core, and enhancing tumor (from large size to small one).

and memory than layerwise NP, these reduced resources in 2D CNNs are much smaller than those in the 3D CNNs while our RANP_f reduces more resources than layerwise NP. Meanwhile, ours can reduce much more resources than

SNIP NP and random NP. Hence, Table 4 indices the generalization of our proposal to both 2D and 3D CNNs.

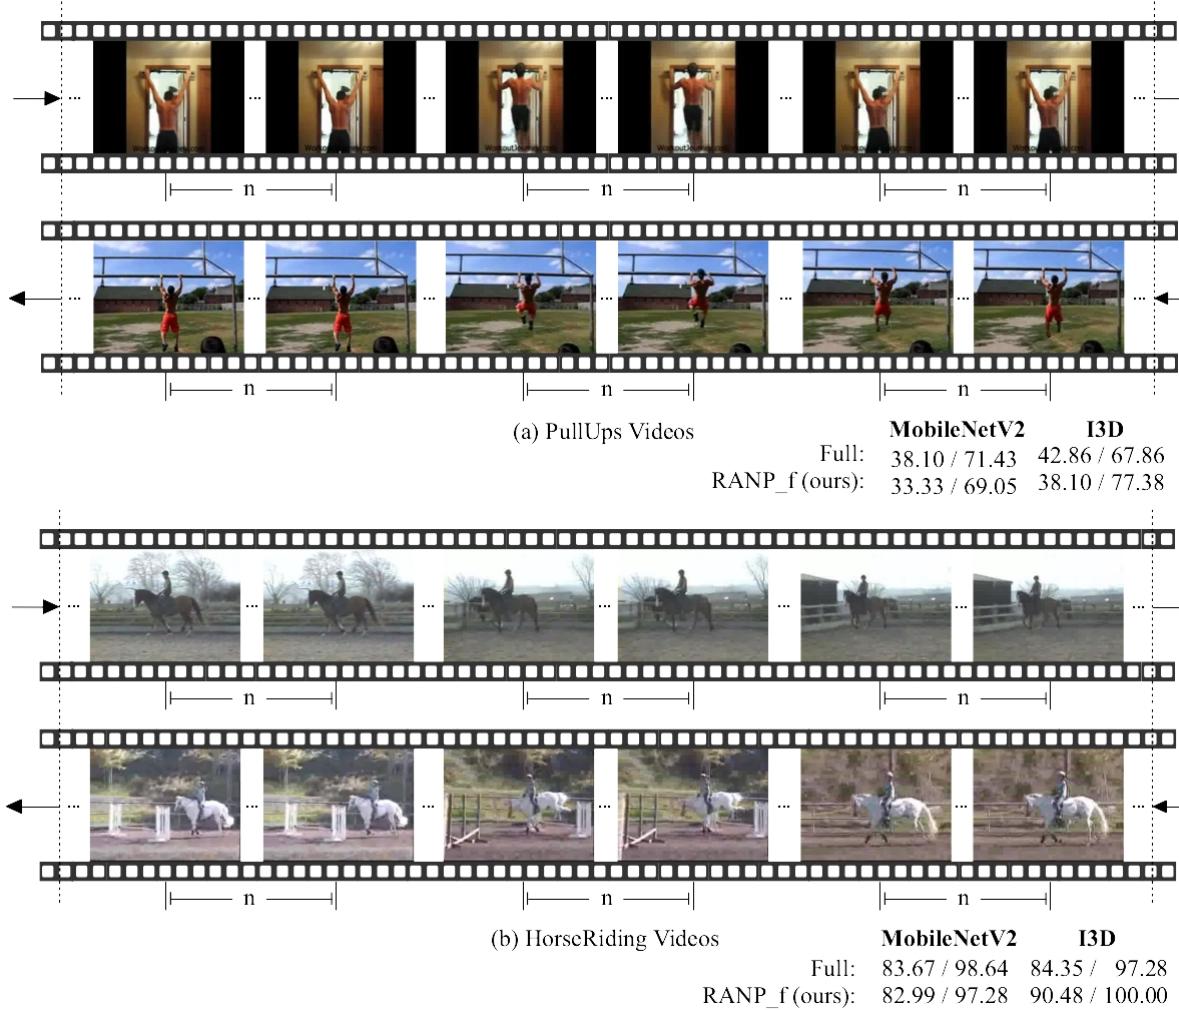


Fig. 7: Examples of UCF101 for video classification using MobileNetV2 and I3D models. Illustrations are on 2 video categories, PullUps and HorseRiding. For every different video, 3 samples with each from $n = 16$ frames are used to calculate accuracy. Metrics “*/” are “Top-1/Top-5”. Pruned networks have different accuracy of different video categories from those of the full networks while the overall accuracy of the whole dataset is in Table 2.

Table 4: Pruning on PSM with 2D-3D CNNs. We mark bold **the best** and underline the second best. Unit of “Param” and “Mem” is MB. Here, 2D and 3D CNN refer to 2D and 3D convolution layers, pooling layers, and batch normalization layers respectively. Resource consumption caused by activation layers is not counted. Overall, RANP_f has a high ability to reduce resources, mainly FLOPs and memory, on both 2D and 3D CNNs without losing accuracy.

Manner	Sparsity	2D CNNs			3D CNNs		
		Param	GFLOPs	Mem	Param	GFLOPs	Mem
Full	0	12.74	174.94	1902.75	7.19	596.08	3595.50
SNIP NP	90.66	8.47	119.89	1674.26	2.63	393.08	3198.23
Random NP		4.38	60.62	1135.57	<u>3.75</u>	310.86	2791.13
Layerwise NP	46.20	3.84	53.06	1051.16	3.87	325.99	2843.16
RANP-f (ours)		4.09	56.28	1079.32	3.83	306.49	2697.61

6.4 Resources and Accuracy with Neuron Sparsity

Here, we further studied the tendencies of resources and accuracy with an increasing neuron sparsity level from 0 to the maximum one with network feasibility.

Resource Reductions. In Figs. 9a, 9c, 9e, 9g, and 9i, RANP, marked with (w), achieves much larger FLOPs and memory reductions than vanilla NP, marked with (w/o), due to the balanced distribution of neuron importance using the reweighting scheme.

Specifically, for *ShapeNet*, RANP prunes up to 98.57% neurons while only up to 78.24% by vanilla NP in Fig. 9a. For *Brats’18*, RANP can prune up to 96.24% neurons while only up to 78.17% neurons can be pruned by vanilla NP in Fig. 9c. For *UCF101*, RANP can prune up to 80.83% neurons compared to 33.15% on MobileNetV2 in Fig. 9e, and

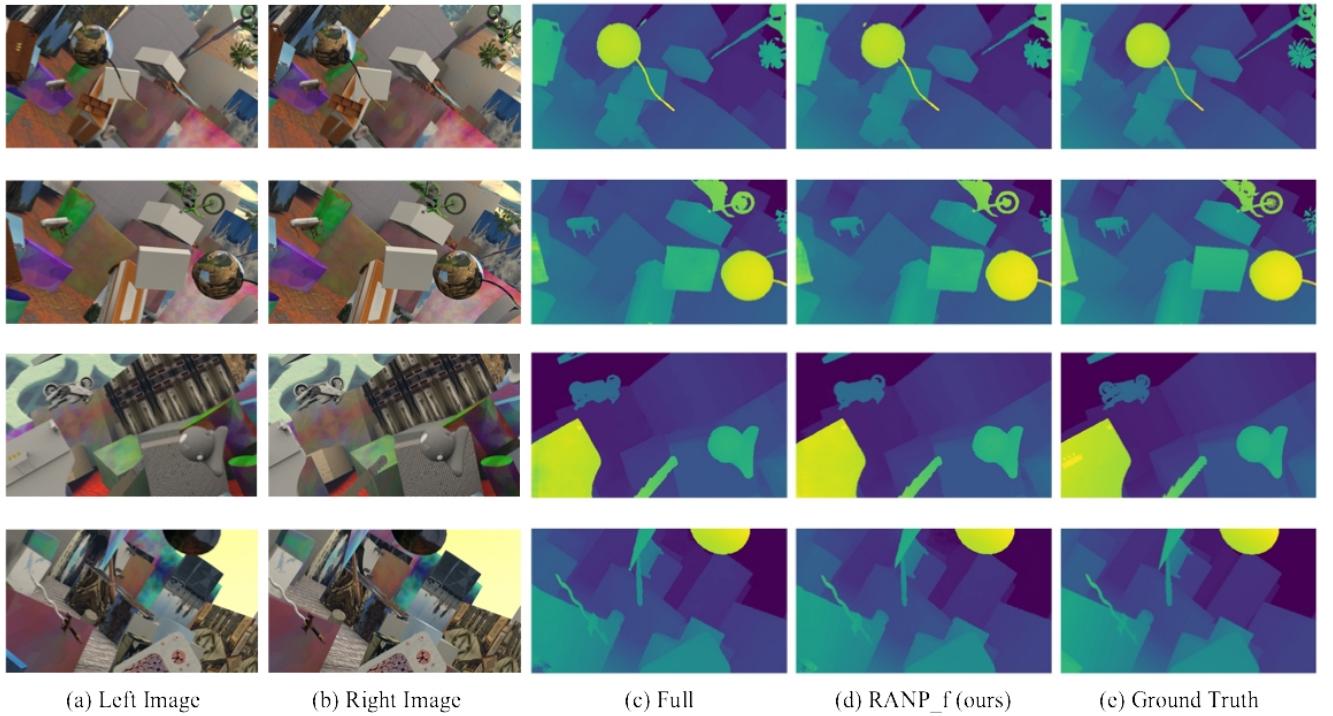


Fig. 8: Examples of SceneFlow for two-view stereo matching using PSM. Each raw is for an example. Predicted disparity maps are in (c)-(d) corresponding to the left images in (a). End-point error and accuracy are shown in Table 2.

85.3% neurons compared to 25.32% on I3D in Fig. 9g. For *SceneFlow* in Fig. 9i, the reweighting scheme induced in RANP_f greatly reduces the resource consumptions with the maximum neuron sparsity 92.32% while only up to 46.20% neuron sparsity can be achieved by vanilla NP.

Accuracy with Pruning Sparsity. For *ShapeNet* in Fig. 9b, the 23-layer 3D-UNet achieves a higher mIoU than the 15-layer one. Extremely, when pruned with the maximum neuron sparsity 97.99%, it can achieve 78.10% mIoU. With the maximum neuron sparsity 98.57%, however, the 15-layer 3D-UNet achieves only 61.42%.

For *BrATS’18* in Fig. 9d, the 23-layer 3D-UNet does not always outperform the 15-layer one and has a larger fluctuation which could be caused by the limited training samples. Nevertheless, even in the extreme case, the 23-layer 3D-UNet has small accuracy loss. Clearly, RANP makes it feasible to use deeper 3D-UNets without the memory issue.

For *UCF101* in Figs. 9f and 9h, RANP-f achieves <3% accuracy loss at 70% sparsity, indicating its effectiveness of greatly reducing resources with small accuracy loss.

For *SceneFlow* in Figs. 9j and 9k, pruning at roughly 70% neuron sparsity achieves 2.8% decrease in Acc-1, 1.27% decrease in Acc-3, and 0.3 pixels decrease in EPE while reducing 68.0% FLOPs and 49.3% memory.

6.5 Initialization for Neuron Imbalance

The imbalanced layer-wise distribution of neuron importance hinders pruning at a high sparsity level due to the pruning of the whole layer(s). For 2D classification tasks in Lee et al. (2020), orthogonal initialization is used to effectively solve this problem for balancing the importance of parameters; but it does not improve our neuron pruning results in 3D tasks and even leads to a poor pruning capability with a lower maximum neuron sparsity than Glorot initialization Glorot and Bengio (2010), which is mentioned in Sec. 4.1. Here, we compare the resource reducing capability using Glorot initialization and orthogonal initialization.

Resource Reductions. In Table 5, vanilla neuron pruning with Glorot initialization, *i.e.*, vanilla-xn, achieves smaller FLOPs and memory consumption than those with orthogonal initialization, *i.e.*, vanilla-ort, except FLOPs with 3D-UNet on *ShapeNet*, I3D on *UCF101*, and PSM on *SceneFlow*. This exception of I3D on *UCF101* is possibly caused by the high ratio of 1^3 kernel size filters in I3D, *i.e.*, 37 out of 57 convolutional layers, because those 1^3 kernel size filters can be regarded as 2D filters on which orthogonal initialization can effectively deal with Lee et al. (2020). While this ratio is also high in MobileNetV2, *i.e.*, 34 out of 52 convolutional layers, it is unnecessary to have the same problem as I3D since it is also affected by the number of neurons in each layer. Note that since 3D-UNets used are all with 3^3

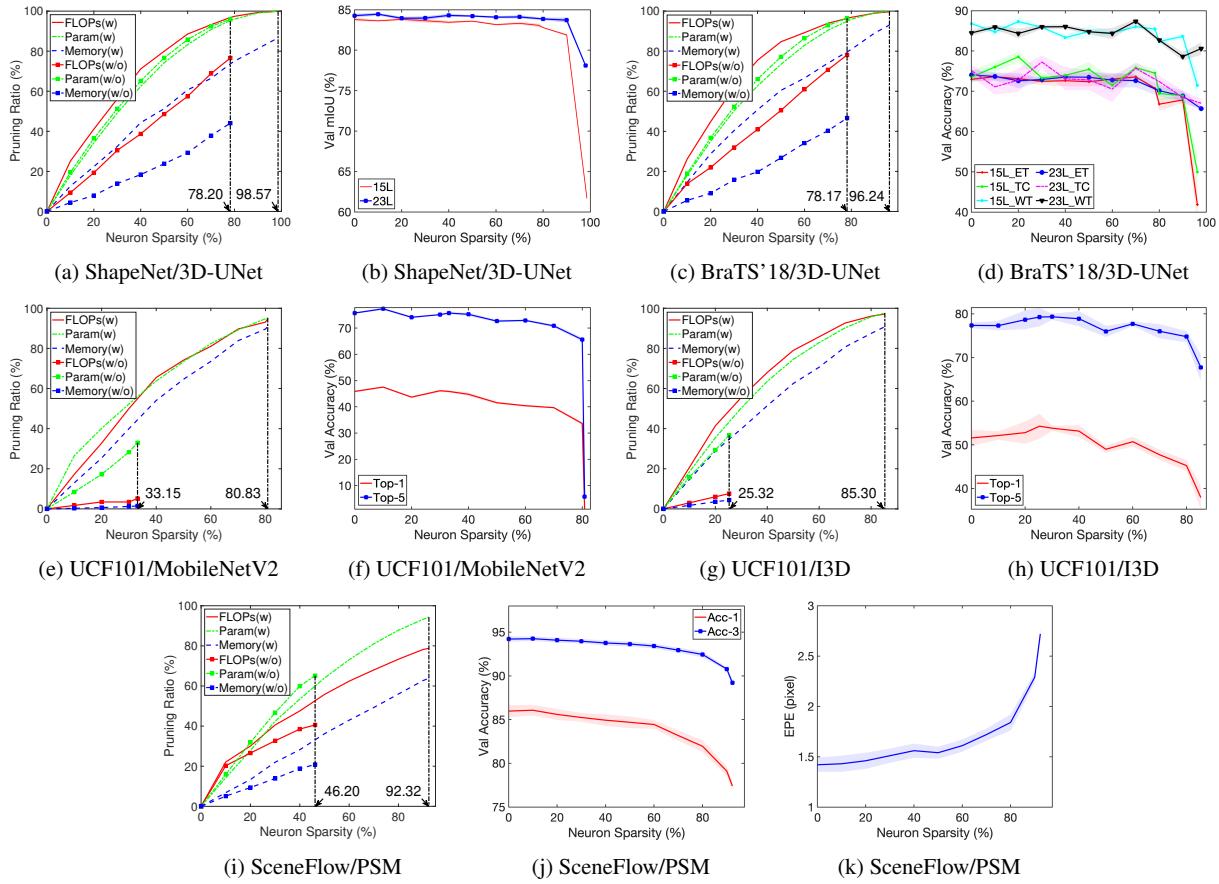


Fig. 9: With minimal accuracy loss, more resources are reduced with (w) reweighting by RANP-f than without (w/o) by vanilla NP. (a), (c), (e), (g), and (i) are resource reductions with sparsity (w) and (w/o) reweighting; the rests are accuracy with sparsity. Best view in color.

kernel size filters, the orthogonal initialization for 3D-UNet in most cases is inferior to Glorot initialization according to our experiments.

Meanwhile, the exception of PSM on SceneFlow is possibly caused by the massive cross-layer additions, which breaks the independence of nonadjacent layers such that retained neurons in each layer have a strong connection with multiple layers, as described in “3D CNNs” in Sec. 6.1.

Furthermore, in Table 5, the gap between vanilla-ort and vanilla-xn is very small on MobileNetV2 and I3D. Nevertheless, with RANP-f and Glorot initialization, *i.e.*, RANP-f-xn, more FLOPs and memory can be reduced than using orthogonal initialization, *i.e.*, RANP-f-ort.

Balance of Neuron Importance Distribution. More importantly, with reweighting by RANP in Fig. 10, the values of neuron importance are more balanced and stable than those of vanilla neuron importance. This can largely avoid network infeasibility without pruning the whole layer(s).

Now, we analyse the neuron distribution from the observation of neuron importance values and network struc-

tures. Fig. 11 illustrates a detailed comparison between orthogonal and Glorot initialization by each two subfigures in column of Fig. 10. In Figs. 11a-11c, vanilla neuron importance by Glorot initialization is more stable and compact than that by orthogonal initialization. After applying the reweighting scheme of RANP-f, the importance tends to be in a similar tendency, shown in Figs. 11b-11d. Consequently, in Figs. 11e-11f, neuron ratios are more balanced after the reweighting than without it, especially the 8th layer. Thus, we choose Glorot initialization as network initialization. Note that we adopt the same neuron sparsity for these two initialization experiments in Table 5 and Fig. 11.

6.6 Transferability with Interactive Model

In this experiment, we trained on ShapeNet with a transferred 3D-UNet by RANP on BraTS’18 with 80% neuron sparsity. Interactively, with the same neuron sparsity, a transferred 3D-UNet by RANP on ShapeNet was applied to train on BraTS’18. Results in Table 6 demonstrate that training

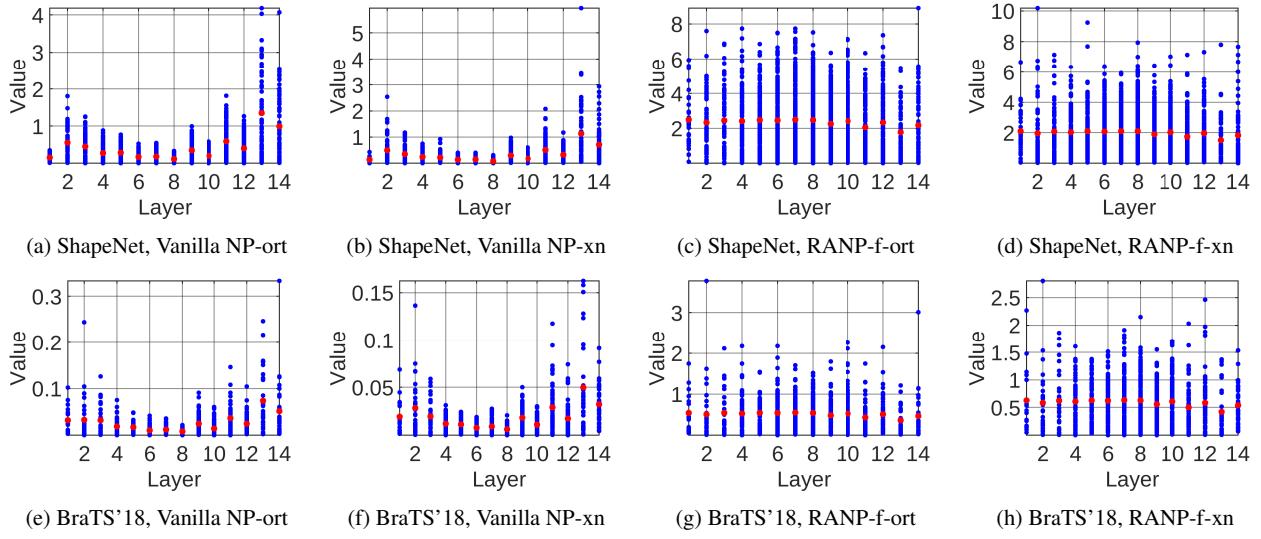


Fig. 10: Neuron importance of 15-layer 3D-UNet by MPMG-sum with orthogonal and Glorot initialization. Blue: neuron values; red: mean values. By vanilla NP, orthogonal initialization does not result in a balanced neuron importance distribution compared to Glorot initialization whereas by our RANP-f, the values are more balanced and resource aware on FLOPs, enabling pruning at the extreme sparsity.

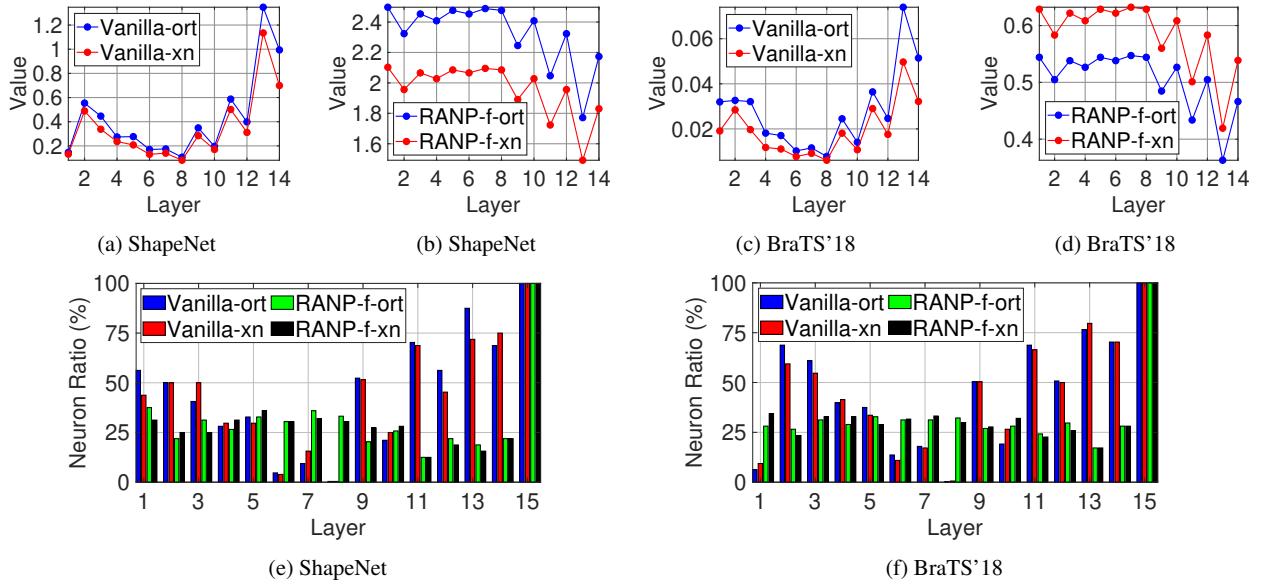


Fig. 11: Comparison of neuron distribution with orthogonal and Glorot initialization before and after reweighting. (a)-(d) are neuron importance values. (e)-(f) are neuron retained ratios. Vanilla versions (both orthogonal and Glorot initializations) prune all the neuron in layer 8, leading to network infeasibility while our RANP-f versions have a balanced distribution of retained neurons.

with transferred models crossing different datasets can largely maintain high or higher accuracy.

6.7 Lightweight Training on a Single GPU

RANP with high neuron sparsity makes it feasible to train with large data size on a single GPU due to the largely reduced resources. We trained on ShapeNet with the same

Table 5: Impact of parameter initialization on neuron pruning. “ort”: orthogonal initialization; “xn”: Glorot initialization; “f”: FLOPs. “Sparsity” is the least max neuron sparsity among all manners to ensure the network feasibility. RANP-f with Glorot initialization achieves the least FLOPs and memory consumption.

Dataset (Model)	Manner	Sparsity (%)	Param (MB)	GFLOPs	Mem (MB)
ShapeNet (3D-UNet)	Full	0	62.26	237.85	997.00
	Vanilla-ort		4.40	72.65	630.00
	Vanilla-xn		4.56	73.22	618.35
	RANP-f-ort	70.53	5.40	21.73	366.29
	RANP-f-xn		5.52	15.06	328.66
BraTS’18 (3D-UNet)	Full	0	15.57	478.13	3628.00
	Vanilla-ort		0.95	159.91	2240.33
	Vanilla-xn		0.92	130.28	2109.19
	RANP-f-ort	72.20	1.24	33.28	967.56
	RANP-f-xn		1.29	23.31	850.56
UCF101 (MobileNetV2)	Full	0	9.47	0.58	157.47
	Vanilla-ort		6.80	0.56	155.71
	Vanilla-xn		6.77	0.55	155.48
	RANP-f-ort	30.21	5.12	0.32	105.88
	RANP-f-xn		5.19	0.28	94.50
UCF101 (I3D)	Full	0	47.27	27.88	201.28
	Vanilla-ort		30.56	25.83	192.70
	Vanilla-xn		30.64	25.85	192.88
	RANP-f-ort	24.24	27.39	15.94	144.10
	RANP-f-xn		27.38	14.63	133.80
SceneFlow (PSM)	Full	0	19.93	771.02	7217.37
	Vanilla-ort		8.66	485.98	5913.52
	Vanilla-xn		8.75	486.92	5947.86
	RANP-f-ort	36.76	10.03	422.34	5332.72
	RANP-f-xn		9.96	420.77	5313.30

Table 6: Transfer learning by 23-layer 3D-UNets interactively pruned and trained between ShapeNet and BraTS’18. Accuracy loss from RANP-f to T-RANP-f is negligible. “T”: transferred.

Manner	ShapeNet		BraTS’18		
	mIoU (%)	ET (%)	TC (%)	WT (%)	
Full	84.27±0.21	74.04±1.45	75.11±2.43	84.49±0.74	
RANP-f (ours)	<u>83.86±0.15</u>	71.13±1.43	72.40±1.48	83.32±0.62	
T-RANP-f (ours)	83.25±0.17	<u>72.74±0.69</u>	<u>73.25±1.69</u>	85.22±0.57	

Table 7: ShapeNet: a deeper 23-layer 3D-UNet is achievable on a single GPU with 80% neuron pruning.

Manner	Layer	Batch	GPU(s)	Sparsity (%)	mIoU (%)
Full	15	12	2	0	83.79±0.21
Full	23	12	2	0	<u>84.27±0.21</u>
RANP-f (ours)	23	12	1	80	84.34±0.21

batch size 12 and spatial size 64^3 in Sec. 6.1 using a 23-layer 3D-UNet with 80% neuron sparsity on a single GPU. With this setup, RANP-f reduces $\sim 35\times$ GFLOPs (from 259.59 to 7.39) and $\sim 3.9\times$ memory (from 1005.96MB to 255.57MB), making it feasible to train on a single GPU instead of 2 GPUs. It achieves a higher mIoU, $84.34\pm0.21\%$, than the 15-layer and 23-layer full 3D-UNets in Table 7.

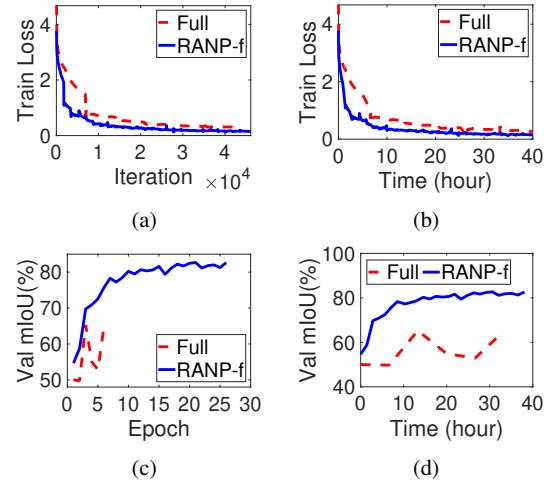


Fig. 12: ShapeNet: a faster convergence on a single GPU with 23-layer 3D-UNet and increased batch size due to the largely reduced resources by RANP-f. Batch size is 1 for “Full” and 4 for “RANP-f”. Experiments run for 40 hours.

The accuracy increase is due to the enlarged batch size on each GPU. With limited memory, however, training a 23-layer full 3D-UNet on a single GPU is infeasible.

6.8 Fast Training with Increased Batch Size

Here, we used the largest spatial size 128^3 of one sample on a single GPU and then extended it to RANP with increased batch size from 1 to 4 to fully fill GPU capacity. The initial learning rate was reduced from 0.1 to 0.01 due to the batch size decreased from 12 in Table 7. This is to avoid an immediate increase in training loss right after 1st epoch due to the unsuitably large learning space.

In Fig. 12a, RANP-f enables increased batch size 4 and achieves a faster loss convergence than the full network. In Fig. 12c, the full network executed 6 epochs while RANP-f reached 26 epochs. Vividly shown by training time in Figs. 12b and 12d, RANP-f has much lower loss and higher accuracy than the full one. This greatly indicates the practical advantage of RANP on fastening training convergence.

7 Conclusion

In this paper, we propose an effective resource aware neuron pruning method, RANP, for 3D CNNs. RANP prunes a network at initialization by greatly reducing resources with negligible loss of accuracy. Its resource aware reweighting scheme balances the neuron importance distribution in each layer and enhances the pruning capability of removing a high ratio, say 80% on 3D-UNet, of neurons with minimal

accuracy loss. This advantage enables training deep 3D CNNs with a large batch size to improve accuracy and achieving lightweight training on one GPU.

Experiments on 3D semantic segmentation using ShapeNet and BraTS'18, video classification using UCF101, and two-view stereo matching using SceneFlow demonstrate the effectiveness of RANP by pruning 70%-80% neurons with minimal loss of accuracy. Moreover, the transferred models pruned on a dataset and trained on another one are succeeded in maintaining high accuracy, indicating the high transferability of RANP. Meanwhile, the largely reduced computational resources enable lightweight and fast training on one GPU with increased batch size.

Acknowledgement

We would like to thank Dr. Ondrej Miksik for valuable discussions. This work is supported by the Australian Centre for Robotic Vision (CE140100016) (Canberra, Australia), Microsoft Research (Redmond, USA), and Data61, CSIRO (Canberra, Australia).

References

- Bakas S, Akbari H, Sotiras A, Bilello M, Rozycski M, Kirby J, Freymann J, Farahani K, Davatzikos C (2017) Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Nature Scientific Data*
- Carreira J, Zisserman A (2017) Quo vadis, action recognition? a new model and the Kinetics dataset. *CVPR*
- Chang J, Chen Y (2018) Pyramid stereo matching network. *CVPR*
- Chen C, Tung F, Vedula N, Mori G (2018) Constraint-aware deep neural network compression. *ECCV*
- Chen H, Wang Y, Shu H, Tang Y, Xu C, Shi B, Xu C, Tian Q, Xu C (2020) Frequency domain compact 3D convolutional neural networks. *CVPR*
- Cicek O, Abdulkadir A, Lienkamp S, Brox T, Ronneberger O (2016) 3D U-Net: Learning dense volumetric segmentation from sparse annotation. *MICCAI*
- Dong X, Chen S, Pan S (2017) Learning to prune deep neural networks via layer-wise optimal brain surgeon. *NeurIPS*
- Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. *International Conference on Artificial Intelligence and Statistics (AISTATS)*
- Gordon A, Eban E, Nachum O, Chen B, Wu H, Yang T, Choi E (2018) Morphnet: Fast & simple resource-constrained structure learning of deep networks. *CVPR*
- Graham B, Engelcke M, Maaten L (2018) 3D semantic segmentation with submanifold sparse convolutional networks. *CVPR*
- Guo Y, Yao A, Chen Y (2016) Dynamic network surgery for efficient dnns. *NeurIPS*
- Han S, Pool J, Tran J, Dally W (2015) Learning both weights and connections for efficient neural network. *NeurIPS*
- He Y, Zhang X, Sun J (2017) Channel pruning for accelerating very deep neural networks. *ICCV*
- He Y, Lin J, Liu Z, Wang H, Li L, Han S (2018) Amc: Automl for model compression and acceleration on mobile devices. *ECCV*
- Hou R, Chen C, Sukthankar R, Shah M (2019) An efficient 3D CNN for action/object segmentation in video. *BMVC*
- Huang Z, Wang N (2018) Data-driven sparse structure selection for deep neural networks. *ECCV*
- Ji S, Xu W, Yang M, Yu K (2013) 3D convolutional neural networks for human action recognition. *TPAMI*
- Kao P, Ngo T, Zhang A, Chen J, Manjunath B (2018) Brain tumor segmentation and tractographic feature extraction from structural MR images for overall survival prediction. *Workshop on MICCAI*
- Karpathy A, Toderici G, Shetty S, Leung T, Sukthankar R, Fei-Fei L (2014) Large-scale video classification with convolutional neural networks. *CVPR*
- Kingma D, Ba J (2015) Adam: A method for stochastic optimization. *ICLR*
- Kleesiek J, Urban G, Hubert A, Schwarz D, Hein K, Bendszus M, Biller A (2016) Deep MRI brain extraction: A 3D convolutional neural network for skull stripping. *NeuroImage*
- Kopuklu O, Kose N, Gunduz A, Rigoll G (2019) Resource efficient 3D convolutional neural networks. *ICCVW*
- Lee N, Ajanthan T, Torr P (2019) SNIP: Single-shot network pruning based on connection sensitivity. *ICLR*
- Lee N, Ajanthan T, Gould S, Torr P (2020) A signal propagation perspective for pruning neural networks at initialization. *ICLR*
- Li C, Wang Z, Wang X, Qi H (2019) Single-shot channel pruning based on alternating direction method of multipliers. *arXiv:190206382*
- Li H, Kadav A, Samet IDH, Graf H (2016) Pruning filters for efficient convnets. *arXiv preprint arXiv:160808710*
- Mayer N, Ilg E, Hausser P, Fischer P, Cremers D, Dosovitskiy A, Brox T (2016) A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. *CVPR*
- Menze B, Jakab A, et al SB (2015) The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*
- Molchanov P, Tyree S, Karras T, Aila T, Kautz J (2017) Pruning convolutional neural networks for resource efficient inference. *ICLR*
- Qi C, Su H, Mo K, Guibas L (2017) Pointnet: Deep learning on point sets for 3D classification and segmentation. *CVPR*
- Reddi S, Kale S, Kumar S (2018) On the convergence of adam and beyond. *ICLR*
- Riegler G, Ulusoy A, Geiger A (2017) OctNet: Learning deep 3D representations at high resolutions. *CVPR*
- Sandler M, Howard A, Zhu M, Zhmoginov A, Chen L (2018) MobileNetv2: Inverted residuals and linear bottlenecks. *CVPR*
- Saxe A, McClelland J, Ganguli S (2014) Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *ICLR*
- Simonyan K, Zisserman A (2014) Two-stream convolutional networks for action recognition in videos. *NeurIPS*
- Soomro K, Zamir A, Shah M (2012) UCF101: A dataset of 101 human action classes from videos in the wild. *CRCV-Techinal Report*
- Sutskever I, Martens J, Dahl G, Hinton G (2013) On the importance of initialization and momentum in deep learning. *ICML*
- Wang C, Zhang G, Grosse R (2020) Picking winning tickets before training by preserving gradient flow. *ICLR*
- Wen W, Wu C, Wang Y, Chen Y, Li H (2016) Learning structured sparsity in deep neural networks. *NeurIPS*
- Yi L, Kim V, Ceylan D, Shen I, Yan M, Su H, Lu C, Huang Q, Sheffer A, Guibas L (2016) A scalable active framework for region annotation in 3D shape collections. *SIGGRAPH Asia*
- Yi L, Shao L, Savva M (2017) Large-scale 3D shape reconstruction and segmentation from shapenet core55. *arXiv preprint arXiv:171006104*
- Yu J, Huang T (2019) Autoslim: Towards one-shot architecture search for channel numbers. *arXiv:190311728*
- Yu N, Qiu S, Hu X, Li J (2017) Accelerating convolutional neural networks by group-wise 2D-filter pruning. *IJCNN*
- Yu R, Li A, Chen C, Lai J, Morariu V, Han X, Gao M, Lin C, Davis L (2018) NISP: Pruning networks using neuron importance score propagation. *CVPR*

- Zanjani F, Moin D, Verheij B, Claessen F, Cherici T, Tan T, With P (2019) Deep learning approach to semantic segmentation in 3D point cloud intra-oral scans of teeth. Proceedings of Machine Learning Research
- Zhang C, Luo W, Urtasun R (2018) Efficient convolutions for real-time semantic segmentation of 3D point cloud. 3DV
- Zhang H, Jiang K, Zhang Y, Li Q, Xia C, Chen X (2014) Discriminative feature learning for video semantic segmentation. International Conference on Virtual Reality and Visualization
- Zhang M, Stadie B (2019) One-shot pruning of recurrent neural networks by jacobian spectrum evaluation. arXiv:191200120
- Zhang Y, Wang H, Luo Y, Yu L, Hu H, Shan H, Quek T (2019) Three dimensional convolutional neural network pruning with regularization-based method. ICIP